How Does Consumption Respond to a Transitory Income Shock? Reconciling Natural Experiments and Semi-Structural Estimations

Jeanne Commault*

October 2018
(Click here for the most recent version)

Abstract

Studies based on natural experiments find that consumption responds strongly and significantly to a transitory variation in income such as a tax rebate or a lottery win. Contrary to this, semi-structural methods that rely on theoretical restrictions to identify the consumption elasticity to a transitory shock in longitudinal survey data find that this elasticity is small and not statistically significant. I show that the two approaches reconcile when relaxing the assumption made by semi-structural methods that the log-consumption growth of a household is independent of the income shocks it has received in the past. With this generalized semi-structural method, the elasticity of nondurable consumption to a transitory shock is 0.54 and statistically significant. It implies that the marginal propensity to consume nondurables out of a transitory income shock over the following year is at least 0.24, which is consistent with findings from natural experiments.

Key words: Marginal propensity to consume, transitory shocks, precautionary behavior, life-cycle model

*Department of Economics, Sciences Po, 28 Rue des Saints-Pères, 75007 Paris, France; jeanne.commault@sciencespo.fr. I thank Richard Blundell, Stéphane Bonhomme, Pierre Cahuc, Edouard Challe, Mariacristina De Nardi, Juan Dolado, Stéphane Dupraz, Jean-Marc Robin, Ben Etheridge, Francis Kramarz, Hamish Low, Eric Mengus, Xavier Ragot, Michèle Tertilt, and participants at seminars where the paper was presented for helpful comments.

Introduction

How does individual consumption respond to a transitory income shock? The answer has implications for a number of economic questions, including the effect of a fiscal shock, the relation between income and consumption inequalities, and the dynamics of business cycles. One obstacle in the way of measuring this response, however, is that transitory shocks are not usually observed directly. Instead, in longitudinal survey data, households report their total income change, without distinguishing between transitory and permanent changes. Yet the difference is important to a researcher, because shocks that do not have the same durability should have distinct effects on consumption.

There are two main solutions to overcome this issue, but they yield opposite conclusions. A first approach consists in exploiting specific episodes of observed transitory income variations, such as a tax rebate and a lottery win, and pairing them with consumption data to directly measure the response of expenditures to an income shock that the researcher observes and knows to be transitory. In the great majority of these studies, a transitory income change has a statistically significant and economically large effect on consumption.[1] A second approach identifies the response of consumption to a transitory shock by putting more structure on the data. Making assumptions about the form of the income process and the way households take their consumption decisions, the seminal paper of Blundell, Pistaferri, and Preston (2008) (hereafter BPP) derives restrictions that can separately identify the elasticity of consumption to a transitory and to a permanent shock. Specifically, the authors assume that income evolves as a transitory-permanent process and that the log-consumption growth of a household does not depend on the income shocks it has received in the past, which is justified as a condition that would hold if the household were to solve the standard life-cycle model.[2] This identification strategy is now influential in all fields of economics that are concerned with the way shocks are passed on to households' decision variables such as consumption and saving but also individual labor supply and time allocation. Yet, the studies relying on

[1]See for example Parker (1999), Souleles (2002) for a transitory change in take-home pay; Souleles (1999), Johnson, Parker, and Souleles (2006), Agarwal, Liu, and Souleles (2007), Parker, Souleles, Johnson, and McClelland (2013), Kaplan and Violante (2014), Misra and Surico (2014) for the consumption response to a tax refund or tax rebate; Baker and Yannelis (2017), Gelman, Kariv, Shapiro, Silverman, and Tadelis (2018) for the response to the 2013 government shutdown; Agarwal and Qian (2014), Kan, Peng, and Wang (2017) for the response to the distribution of cash or consumption vouchers by the government; Fagereng, Holm, and Natvik (2018) for the response to a lottery win. A related but distinct literature relies on hypothetical survey responses rather than direct observations of consumption to measure how households respond to a transitory shock. See Parker and Souleles (2017) for a comparison between hypothetical survey measures and natural experiment measures.

[2]This method is said to be semi-structural because it does not require that the fully-fledged life-cycle model holds but only that this one condition derived from it does, together with the transitory-permanent specification of income.

this estimation method typically find that the elasticity of consumption to a transitory shock is not statistically significant, although it is quite precisely estimated.[3] The divergence between the two approaches means that either the particular transitory shocks observed in natural experiments share specific characteristics inducing households to respond more to them than they do to the shocks they face in survey data, which would suggest that these natural experiments have little external validity, or that some of the identifying restrictions imposed by semi-structural methods do not hold well enough in the data to yield reliable estimates.

In this paper, I generalize the semi-structural estimation method by relaxing the assumption that the log-consumption growth of a household does not depend on the income shocks it has received in the past, and I find that the elasticity of consumption to a transitory shock is statistically significant and within the range of values suggested by natural experiments. More precisely, I make three contributions: (i) I motivate my generalization by proving that log-consumption growth depends on the realizations of past shocks in the standard life-cycle model, because of precautionary behavior, and I explain that random walk approximations of consumption neglect this effect because they implicitly linearize an identity; (ii) I develop an estimator of the elasticity of consumption to a transitory shock that is robust to the presence of a correlation between log-consumption growth and past transitory shocks, and I show that the BPP estimator of this elasticity is not robust and biased downwards when the correlation is negative; (iii) I implement the robust estimator in the same survey data as BPP and obtain a larger and statistically significant estimate.

As to my first contribution, the correlation between log-consumption growth and past transitory shocks whose existence in the standard life-cycle model I establish is driven by precautionary behavior, that is, the effect of uncertainty on the decisions of the household. Indeed, when utility is isoelastic, marginal utility is convex, so the presence of uncertainty increases the expected marginal utility of future consumption and induces the household to shift some of its resources from the present to the future, raising its consumption growth. The magnitude of this precautionary consumption growth depends on the distribution of future consumption, which is influenced by the level of assets that the household has, thus by the income shocks it has received in the past. Extending the model to incorporate stricter borrowing limits, durable goods, or habit persistence yields additional sources of correlation between log-consumption growth

[3] See section A of the online appendix for a review of the literature that builds on the BPP identification strategy. Early papers before BPP study the response of consumption to a change in total income (Krueger and Perri (2005) Krueger and Perri (2011)), or use specific proxies for permanent and transitory income changes such as disability or short unemployment spells making them close to natural experiments studies (Cochrane (1991), Dynarski and Gruber (1997)).

and past income shocks.

I explain why precautionary consumption growth does not appear in previous approximations of consumption and log-consumption growth, although it would not disappear in first order approximation around small income or consumption innovations because such innovations only take place after the precautionary saving decision has been made. Starting with the seminal expression derived by Hall (1978), random walk approximations are based on the linearization of an identity rather than on the linearization of the first order condition of the household maximization problem, because this condition does not relates current consumption with realized future consumption but only with the expected distribution of future consumption. Trying to linearize realized future consumption from it leads to linearizing a relation between future consumption and itself. The random walk expression obtained is contingent upon the choice of the identity that is linearized. The approximation of log-consumption growth on which BPP rely is obtained from a similar procedure as the one developed by Hall (1978).

For my second contribution, I build a robust estimator of the elasticity of consumption to a transitory shock that is valid in the presence of a correlation with past shocks. The statistical model on which the estimator is based encompasses the standard life-cycle model as a special case, allowing for a broader range of behaviors. In particular, it does not require that a household decides on its consumption by solving a maximization problem. To identify the effect of an unobserved transitory shock on log-consumption growth, the strategy is to use future log-income growth as an instrument. Indeed, when log-income evolves as a transitory-permanent process, future log-income growth correlates negatively with the current transitory shock, since the effect of the transitory shock fades away, but is independent of the current permanent shock. In addition, future log-income growth at the last period in the future before the effect of a current transitory shock fully dissipates correlates with the current transitory shock but not with past transitory shocks, whose effect is already dissipated.

However, neglecting the correlation between log-consumption growth and past transitory shocks leads a researcher to use future log-income growth at all periods as an instrument, and not only future log-income growth at the last period before the current transitory shock fully dissipates. Such a set of instruments does not separately identify the positive effect of the current transitory shock from the negative effect of the past transitory shocks on log-consumption growth, so the response of consumption to the current transitory shock is underestimated.

Regarding my third contribution, I implement the robust estimator in data from the Panel Study of Income Dynamics (PSID) between 1978 and 1992, combined with consumption data imputed from the Consumer Expenditure Survey (CEX) over the same period. This follows BPP, making the comparison between my results and theirs straightforward. With the robust estimator, the average elasticity of nondurable consumption to a transitory shock over the following year has a point estimate of 0.54 and is statistically significant. Contrary to this, an estimator that includes future income growth at periods when it correlates with past transitory shocks in the set of instruments yields a point estimate of 0.05, not statistically significant. The BPP estimator, which also uses future income growth at periods when it correlates with past transitory shocks as instrument but additionally relies on other moments to estimate extra parameters, obtains a similar point estimate of 0.05, not statistically significant. The robust elasticity estimate of 0.54 implies an average marginal propensity to consume (MPC) of at least 0.24, which is consistent with results from natural experiments. Breaking down the sample in subgroups, the elasticity is less precisely estimated but the point estimate is higher for households with lower financial income, a lower wage rate, and a younger head.

I consider three extensions in which: (i) permanent income is an AR(1), that is, the impact of a permanent shock can depreciate over time, (ii) the households may partly anticipate the value of their future transitory and permanent shocks, and (iii) measurement error can be serially correlated. My finding that the elasticity of consumption to a transitory shock is large and significant is unaffected by these additional features.

Considering semi-structural estimation techniques that are related, I show that estimators of the elasticity of consumption to a permanent are not robust to the presence of a correlation between log-consumption growth and past shocks either. Also, pioneered by Blundell, Pistaferri, and Saporta-Eksten (2016), a strand of the literature goes one step further and aims to measure the Frisch elasticities of the households, that is, their elasticities at constant marginal utility of wealth, to disentangle the wealth effect of a shock from the effect of the adjustments in other margins that it may trigger. These Frisch elasticities are estimated by matching the empirical values of the standard elasticities, estimated with a BPP method, with their theoretical expressions as functions of the Frisch elasticities. This poses two problems when households have a precautionary motive and their log-consumption growth correlates with their past shocks: first, the empirical values estimated with a BPP method are biased; second, the theoretical expressions are approximations that neglect the contribution of precautionary behavior, inducing further bias. Finally, relying on biennial data when income shocks take place annually also induces a bias in the measured elasticity of consumption to a transitory shock, because it neglects the effect of the transitory shock occurring in between the two years, which can correlate with log-consumption growth.

Two papers in the literature closely relate to mine. Kaplan and Violante (2010) make two early points on the validity of the BPP estimator: (i) that it is uncertain whether the assumption of no correlation with past shocks holds in the standard incomplete-markets life-cycle model[4]; (ii) that this assumption is not required for estimating the elasticity of consumption to a transitory shock, but only for estimating of the elasticity of consumption to a permanent shock. On simulated data, they test the validity of a version of the BPP estimator that is based on a slightly simpler income process than the one that fits the data used by BPP, which is such that past transitory shocks do not affect the instrument used to identify the elasticity to a transitory shock;[5] they find that this version of the BPP estimator of the elasticity to a transitory shock is unbiased. I further their point (i) by proving analytically that the assumption of no correlation does not hold in the standard life-cycle model, because of precautionary behavior. I further their point (ii) by noting that, although the assumption of no correlation with past shocks is not required to estimate the elasticity of consumption to a transitory shock alone, it is nevertheless imposed by the BPP estimator, and leads to a downward bias in their measure of the elasticity to a transitory shock.

The second paper that I build on is that of Arellano, Blundell, and Bonhomme (2017). The authors generalize the BPP estimator by allowing permanent shocks to be persistent and their persistence to be history-dependent. They also measure the elasticity of consumption conditionally on the asset holdings of the households and on their permanent income. In a world with only a single type of assets, this makes their estimator partly robust to the correlation between log-consumption growth and past transitory shocks caused by precautionary behavior, because most of this correlation plays out through the level of assets.[6] In a world with different types of assets, one would need to control for each type of assets separately. Since the focus of their paper is on the elasticity to a permanent shock, they do not estimate the elasticity to a transitory shock in survey data but only in data simulated from a life-cycle model, and they find no significant response of consumption (Figure S21-S23 in their Supplementary Appendix). Relative to their paper, a first contribution of my paper is that my estimator is robust

[4]See p60 "With respect to assumption (SM) [short memory e.g. no correlation with past permanent shocks one period ago and past transitory shocks two periods ago], one can verify whether it holds in general only in models where the consumption allocation has a closed form. In the absence of a closed form, as in the standard incomplete-markets economy that we study in this paper, one must rely on model simulations.".

[5]They consider an income process in which the transitory component of the income process is an MA(0) and not an MA(1). In that case, there is only one instrument available to identify the elasticity of consumption to a transitory shock, log-income growth at $t+1$, which is log-income growth at a period when the effect of past shocks has already dissipated.: it depends on the realization of the transitory shock at t but not at $t-1$ since transitory shocks have a shorter persistence.

[6]It would not fully control for it, though, because a transitory shock at the immediately past period affects precautionary behavior through its effect on both current assets and current transitory income.

to all types of correlations with past shocks, and not only robust to the correlations in which past income shocks affect current log-consumption growth through their effect on total asset holdings or on permanent income. A second contribution is to implement a robust estimator, not in simulated data but in survey data, in which I find that controlling for the correlation with past shocks makes a difference. A third contribution is that my method does not necessitate observing the level of assets of the households, thus restricting the definition of assets to what is observable.[7]

1 Theoretical motivation

1.1 The standard life-cycle model

Income process Time is discrete and indexed by t = 0, 1, ... , T. The net income of a household i at period t, denoted $y_{i,t}$, is modeled as a transitory-permanent process:

$$ln(y_{i,t}) = p_{i,t} + \mu_{i,t} + f_i + \kappa_t z_{i,t} \tag{2.1}$$
$$with \begin{cases} p_{i,t} & = p_{i,t-1} + \eta_{i,t} \\ \mu_{i,t} & = \varepsilon_{i,t} + \theta_1 \varepsilon_{i,t-1} + ... + \theta_q \varepsilon_{i,t-q}. \end{cases}$$

The log of net income is the sum of a permanent component $p_{i,t}$ that follows a random walk, of a transitory component $\mu_{i,t}$ that follows an MA(q) process, and of a term $f_i + \kappa_t z_{i,t}$ that captures individual fixed effects and the deterministic influence of current demographic characteristics $z_{i,t}$. The uncertainty of the household about its future income comes from the presence of the shocks, $\eta_{i,t}$ and $\varepsilon_{i,t}$. The shock $\eta_{i,t}$ is a permanent shock because its realization remains in the value of p at all following periods, so it affects the income received by the household for the rest of its lifetime. The shock $\varepsilon_{i,t}$ is transitory because its effect on income dissipates after q periods.[8] At each period, the permanent and transitory shocks are drawn independently of each other and independently of their previous realizations. The demographic characteristics $z_{i,t}$ are not subject to any uncertainty: they may change over time, but these variations are expected by the household. Their impact on log-income is measured by the vector of coefficients κ_t, which is allowed to change with calendar time.

[7] In particular, my strategy remains valid if households take into account their expected but not-yet-inherited wealth in their decisions, or rely on other instruments than cash, stocks, bonds, and housing as store of value, which would be unobserved in the data of Arellano, Blundell, and Bonhomme (2017).

[8] By construction, at the end of the household's life the transitory shock resembles a permanent shock (whose effect is depreciating), because its effect might last until the last period of the household's life.

Household's problem At each period t, a household i chooses its current consumption and the distribution of its future consumption as the solution of the following intertemporal optimization problem:

$$\max_{c_{i,t},\dots c_{i,T}} \sum_{s=0}^{T-t} \beta^{t+s} e^{\delta_{t+s} z_{i,t+s}} E_t \left[u(c_{i,t+s}) \right] \tag{2.2}$$

$$s.t. \quad a_{i,t+k+1} = (1+r) a_{i,t+k} - c_{i,t+k} + y_{i,t+k} \qquad \forall\, 0 \leq k \leq T-t, \tag{2.3}$$

$$a_{i,T} \geq 0. \tag{2.4}$$

The household is finite-lived with T the length of its life. It has time-separable preferences, and at each period t it derives utility from its contemporaneous consumption expenditures $c_{i,t}$. The period utility function $u(c)$ is isoelastic so its functional form is $u(c) = \frac{c^{1-\rho}-1}{1-\rho}$. Future utility is discounted by the factor β, and shifted by the demographic characteristics $z_{i,t}$, whose current and future values are known in advance with certainty by the household. The impact of these demographics on utility is captured by the vector of coefficients δ_t, which can change with calendar time. At each period, the household earns the stochastic amount $y_{i,t}$. The budget constraint (2.3) states that to store its wealth from one period to another the household only has access to a risk-free asset that delivers the certain interest rate r, where $a_{i,t}$ denotes the level of this asset at the beginning of period t or at the end of period $t-1$. The terminal condition on wealth (2.4) states that the household cannot die with a strictly positive level of debt.

1.2 Precautionary behavior and correlation with past shocks

First order condition The first order condition of the maximization problem of the household, known as the Euler equation, is as follows:

$$u'(c_{i,t}) = E_t[u'(c_{i,t+1})] R_{i,t,t+1},$$

with $R_{i,t,t+1} = \beta(1+r)e^{\Delta\delta_{t+1} z_{i,t+1}}$ a factor accounting for the deterministic intertemporal substitution motives. It states that an optimizing household chooses its current and future consumption so that they deliver the same expected marginal utility. The natural borrowing limit never binds, as the household would never put itself in the situation of possibly consuming zero in the future by borrowing more than the lowest possible amount that it could earn in the future. The effect of intertemporal substitution can equivalently be expressed as a weight $R_{i,t,t+1}^{1/\rho}$ on current consumption: $u'(c_{i,t}) R_{i,t,t+1}^{-1} = c_{i,t}^{-\rho} R_{i,t,t+1}^{-1} = (c_{i,t} R_{i,t,t+1}^{1/\rho})^{-\rho}$.

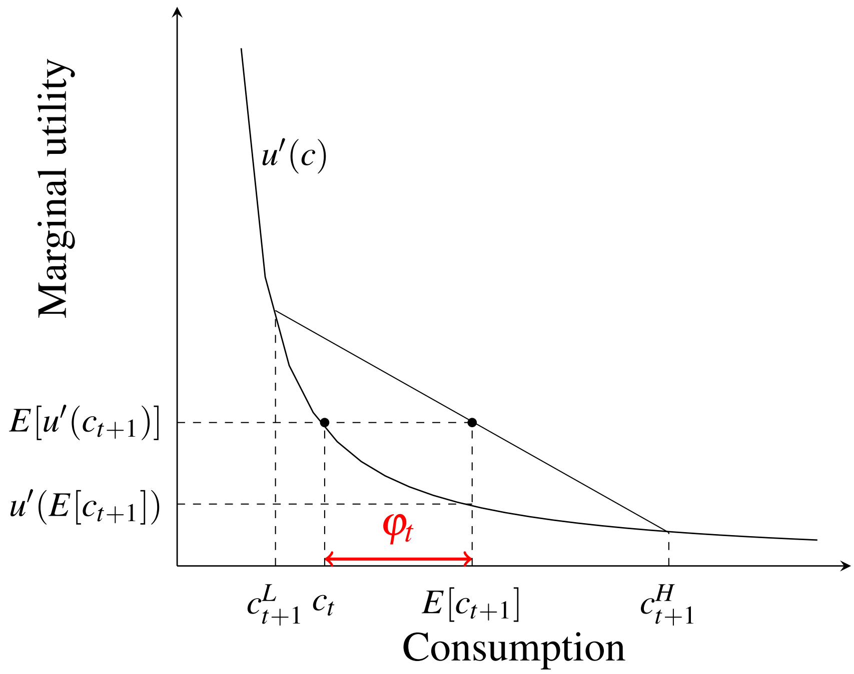

Figure 1: Equalization of expected marginal utility

Note: This graph is illustrative.

Consumption growth When marginal utility is convex, the effects of negative and positive shocks are asymmetric: a negative shock to future consumption raises the value of one additional unit of future consumption more than a positive shock of the same magnitude reduces the value of one additional unit of future consumption. On average, the effect of the negative shocks dominates, so the presence of mean-zero shocks increases the expected marginal utility of future consumption above the marginal utility of expected future consumption: $E_t[u'(c_{i,t+1})] > u'(E_t[c_{i,t+1}])$. This induces the household to set its current consumption below its expected future consumption, so the marginal utility of its current consumption, which is not subject to uncertainty, be as high as the expected marginal utility of its future consumption, which is increased by the uncertainty about future consumption. The amount by which the household must decrease its current consumption coincides with the risk-premium associated with the marginal utility of consumption, that is, the variable $\varphi_{i,t}$ such that: $E_t[u'(c_{i,t+1})] = u'(E_t[c_{i,t+1}] - \varphi_{i,t})$, which is strictly positive when marginal utility is convex:

$$u'(c_{i,t}R^{1/\rho}_{i,t,t+1}) = u'(E_t[c_{i,t+1}] - \varphi_{i,t}),$$
$$E_t[c_{i,t+1}] = c_{i,t}R^{1/\rho}_{i,t,t+1} + \underbrace{\varphi_{i,t}}_{\text{precaution}}.$$

I refer to $\varphi_{i,t}$ as precautionary consumption growth because it corresponds to the effect of uncertainty on consumption growth. Figure 1 presents a graphical illustration of the mechanism in a simple setting. Future consumption can take two values, a low value c_{t+1}^{L} associated with a low income realization, and a high value c_{t+1}^{H} associated with a high income realization. Because marginal utility is convex, a low realization of future consumption increases its marginal utility more than a high realization of future consumption decreases it: the difference between $u'(E_t[c_{t+1}])$ and $u'(c_{t+1}^{L})$ is larger than the difference between $u'(E_t[c_{t+1}])$ and $u'(c_{t+1}^{H})$. Taking their average, the effect of the low realization dominates, raising the expected marginal utility of future consumption above the marginal utility of expected future consumption: $E_t[u'(c_{t+1})]$ is above $u'(E_t[c_{t+1}])$. Current consumption c_t is then chosen as the certain amount of consumption associated with a marginal utility equal to $E_t[u'(c_{t+1})]$. Since $E_t[u'(c_{t+1})]$ is above $u'(E_t[c_{t+1}])$ and marginal utility decreasing in consumption, c_t must be below $E_t[c_{t+1}]$: a household must consume less at the current period than it expects to consume in the future for its current marginal utility to be as large as its expected future marginal utility. There is precautionary consumption growth $\varphi_t > 0$.

Theorem: In the model presented above, the precautionary consumption growth φ_t is negatively correlated with assets and with the value of a past transitory income shock. At any period $0 < t < T$:

$$\frac{d\varphi_{i,t}}{da_{i,t}} < 0 \text{ and } \frac{d\varphi_{i,t}}{d\varepsilon_{i,t-k}} < 0, \quad k \geq q.$$

Proof: The formal derivation is presented in Appendix A. Intuitively, precautionary consumption growth depends on the expected distribution of future consumption, and a gain in assets has two effects on this distribution: (i) it increases future consumption in all states of the world, because it increases the resources available for consumption, shifting the distribution of future consumption upwards to a region over which marginal utility is less convex; (ii) it reduces the variance of future consumption, because having more assets reduces the sensitivity of future consumption to future income shocks, reducing the difference between low and high levels of future consumption. As both effects lessen the extent to which income uncertainty affects the expected marginal utility of future consumption, a gain in assets reduces the need for precautionary consumption growth. Figures 2(a) and 2(b) provide graphical representations. The top figure 2(a) represents effect (i). Keeping constant the difference between the low and high realizations of future consumption, precautionary consumption growth φ_t decreases when the distribution of future consumption shifts upwards, because this moves it to a region

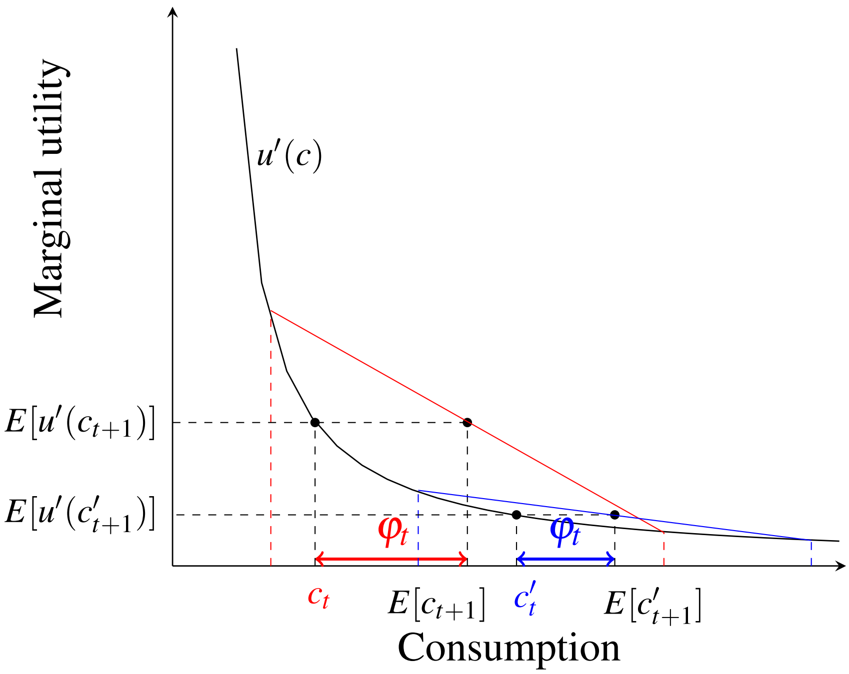

(a) Precautionary consumption growth before (in red) and after (in blue) a shift upwards in the distribution of future consumption

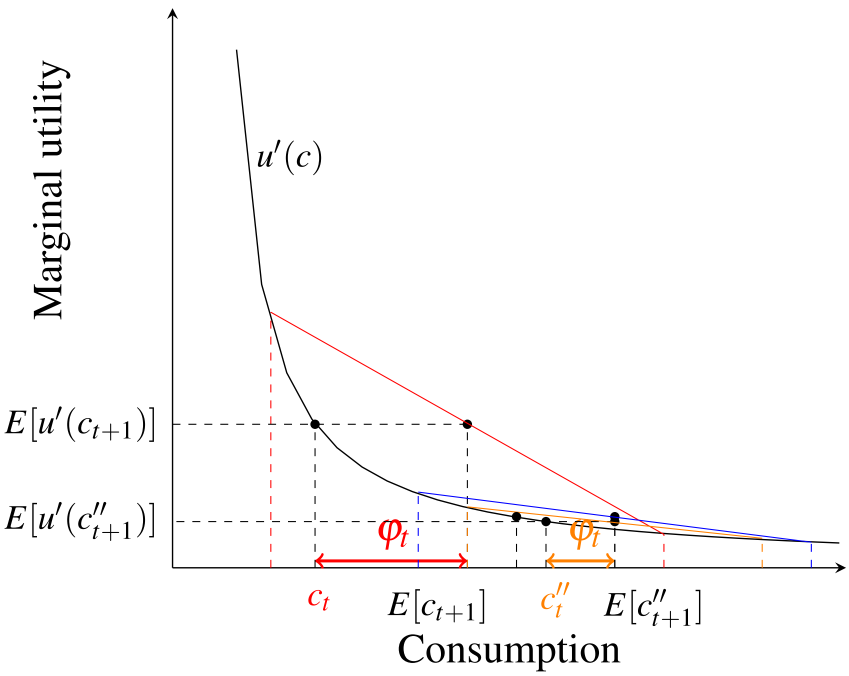

(b) Precautionary consumption growth before (in blue) and after (in orange) a decrease in the variance of future consumption

Figure 2: Equalization of marginal utility before and after a gain in asset

Note: These graphs are illustrative.

over which marginal utility is less convex, thus where the same mean-zero shocks to future consumption have a smaller impact on the expected marginal utility of future consumption. The bottom figure 2(b) pictures effect (ii). Keeping the average value of future consumption constant, precautionary consumption growth φ_t decreases when the difference between the low and high realizations is reduced, because it lowers the magnitude of the mean-zero shocks to future consumption.

Past transitory shocks $\varepsilon_{i,t-k}$, with $k \geq q$ the persistence of transitory income, have the same effect on consumption as assets, because they only influence consumption through their effect on cash-in-hand, and do not affect expected future income.[9] Thus, my theoretical contribution is to establish that the additional consumption growth resulting from precautionary behavior, a phenomenon identified since the '70s,[10] correlates negatively with the level of assets of a household, thus with the transitory income shocks it has received in the past.

Log-consumption growth As precautionary behavior modifies consumption growth, it modifies log-consumption growth as well: the latter incorporates a term that depends positively on precautionary consumption growth φ_t and on the innovation to consumption $\nu_{i,t+1}$, scaled by current consumption $c_{i,t}R_{i,t,t+1}^{1/\rho}$:

$$\Delta ln(c_{i,t+1}) = \underbrace{\frac{1}{\rho} ln(R_{i,t,t+1})}_{\substack{\text{change in dem.} \\ \text{+ int. substitution} \\ \text{(deterministic)}}} + \underbrace{E_t[ln(1+\frac{\varphi_{i,t}+\nu_{i,t+1}}{c_{i,t}R_{i,t,t+1}^{1/\rho}})]}_{\substack{\text{precaution} \\ \text{(correlates with past shocks)}}} + \underbrace{\xi_{i,t+1}}_{\substack{\text{innovation} \\ \text{(uncorrelated} \\ \text{with past shocks)}}} .$$

I refer to the second term on the right-hand-side as the contribution of precaution to expected log-consumption growth, because it would be zero under perfect foresight (i.e. no uncertainty about the future, everything else being equal).[11] Owning more assets reduces $\varphi_{i,t}$, as I establish in the theorem, and increases $c_{i,t}$ so it reduces the ratio $\frac{\varphi_{i,t}}{c_{i,t}R_{i,t,t+1}^{1/\rho}}$. This means that, through this term at least, the log-consumption growth of a household correlates negatively with its level of assets, thus with the value of the transitory shocks

[9]The result does not necessarily hold for $\varepsilon_{i,t-k}$ with $k < q$, because a change in the realization of these shocks does not only change current cash-in-hand but also modifies the expected distribution of future income.

[10]See Leland (1968), Sandmo (1970), Drèze and Modigliani (1972), Sibley (1975), and Miller (1976) for the progressive generalization of this result; Kimball (1990), often cited on the subject, proves in fact the slightly different result that to chose the same level of consumption as it would in the absence of uncertainty a household facing uncertainty must be holding strictly more assets.

[11]This term forms the contribution of precaution to expected log-consumption growth but is not the sole precautionary component of total log-consumption growth, since the distribution of the innovation is affected by precautionary behavior.

it has received in the past.

Common extensions as additional sources of correlation Many extensions of the life-cycle model feature other channels through which past transitory shocks would affect log-consumption growth. I show in section B of the online appendix that the presence of tighter borrowing constraints, the durability of some consumption expenditures, and the formation of habits generate new components of log-consumption growth that correlate with past transitory shocks.

1.3 Pitfalls in approximations of consumption growth

Random walk expression of consumption growth I note that, starting from Hall (1978), random walk expressions of consumption are not based on approximating the first order condition of the household's problem, but on approximating an identity plugged in this first order condition. Hall (1978) starts from the first order condition, substitutes for $E_t[u'(c_{i,t+1})] = u'(c_{i,t+1}) + (E_t[u'(c_{i,t+1})] - u'(c_{i,t+1}))$, and applies $(u')^{-1}(.)$ to each side:

$$\begin{aligned} u'(c_{i,t}R_{i,t,t+1}^{1/\rho}) &= E_t[u'(c_{i,t+1})] \\ u'(c_{i,t}R_{i,t,t+1}^{1/\rho}) &= u'(c_{i,t+1}) + (E_t[u'(c_{i,t+1})] - u'(c_{i,t+1})) \\ c_{i,t+1} &= (u')^{-1}(u'(c_{i,t+1}) + \cancel{E_t[u'(c_{i,t+1})]} - \cancel{u'(c_{i,t}R_{i,t,t+1}^{1/\rho})}). \end{aligned}$$

The crossing of the terms that cancel out is mine. It shows that the Euler equation is unessential here, and that the expression of $c_{i,t+1}$ obtained is the same as the identity $c_{i,t+1} = (u')^{-1}(u'(c_{i,t+1}))$. Indeed, because the first order condition does not establish a relation between current consumption and realized future consumption but only between current consumption and the expected distribution of future consumption, trying to express realized future consumption from it leads to relying on an identity. Hall (1978) then takes a first order approximation of $c_{i,t+1}$ around the point where $u'(c_{i,t+1}) = u'(c_{i,t}R_{i,t,t+1}^{1/\rho})$, i.e. uses $f(x) \approx f(x_0) + (x - x_0)f'(x_0)$ with $f(x) = c_{i,t+1}$, $x = u'(c_{i,t+1})$, and $x_0 = u'(c_{i,t}R_{i,t,t+1}^{1/\rho})$:

$$c_{i,t+1} \approx c_{i,t}R_{i,t,t+1}^{1/\rho} + \underbrace{\frac{u'(c_{i,t+1}) - u'(c_{i,t}R_{i,t,t+1}^{1/\rho})}{u''(c_{i,t}R_{i,t,t+1}^{1/\rho})}}_{\text{uncorrelated with past shocks}}.$$

This is Hall (1978)'s random walk expression, in which future consumption is the sum of current consumption plus a mean zero innovation, uncorrelated with past shocks because $u'(c_{i,t+1}) - u'(c_{i,t}R_{i,t,t+1}^{1/\rho}) = u'(c_{i,t+1}) - E_t[u'(c_{i,t+1})]$. Yet, this expression reflects the choice of the identity that is linearized: if instead one were to approximate $c_{i,t+1} = v^{-1}(v(c_{i,t+1}))$, with $v(.)$ such that $E_t[v(c_{i,t+1})] \neq v(c_{i,t}R_{i,t,t+1}^{1/\rho})$, at the same point as Hall (1978) where $c_{i,t+1} = c_t R_{i,t,t+1}^{1/\rho}$, the resulting expression would not be a random walk, although the first order approximation would hold.[12] The particular identity $c_{t+1} = (u')^{-1}(u'(c_{i,t+1}))$ yields a random walk because it relies on a choice of $x = u'(c_{i,t+1})$ such that $E_t[x] = x_0$.

Although it is not what Hall (1978) does, it would be possible to derive a random walk expression of consumption from the first order condition of the household maximization problem by taking an approximation of current consumption around the point where the expected distribution of future consumption is a Dirac delta function with mass one in current consumption, that is, where future consumption equals current consumption in all states of the world, regardless of the shocks that realize.[13] This is substantially more restrictive than an approximation around the point where realized future consumption equals current consumption, which can occur in the standard life-cycle model with uncertainty, while the point where future consumption equals current consumption in all states of the world cannot emerge as an outcome of the model.

The log-linearized Euler equation is typically obtained with same method as Hall (1978)'s, by linearizing the identity $ln(c_{i,t+1}) = ln(c_{i,t}R_{i,t,t+1}^{1/\rho}) - \frac{1}{\rho} ln(1 + \frac{u'(c_{i,t+1}) - u'(c_{i,t}R_{i,t,t+1}^{1/\rho})}{u'(c_{i,t}R_{i,t,t+1}^{1/\rho})})$ around $u'(c_{i,t+1}) = u'(c_{i,t}R_{i,t,t+1}^{1/\rho})$.

BPP's expression of log-consumption growth The expression used by BPP is based on the one derived by Blundell, Low, and Preston (2013), which partly follows a similar procedure as Hall (1978)'s, applied to log-consumption. I detail formally their steps in section C of the online appendix. The authors obtain that log-consumption growth is uncorrelated with past shocks since it is the sum of a deterministic trend and an

[12]The resulting expression would be $c_{i,t+1} \approx c_{i,t}R_{i,t,t+1}^{1/\rho} + \frac{v(c_{i,t+1}) - v(c_{i,t}R_{i,t,t+1}^{1/\rho})}{v''(c_{i,t}R_{i,t,t+1}^{1/\rho})}$ with $v(c_{i,t}R_{i,t,t+1}^{1/\rho}) \neq E_t[v(c_{i,t+1})]$.

[13]In the discrete case with J possible future states of the world, and each state $j \in J$ occurring with probability π_j:

$$u'(c_{i,t}R_{i,t,t+1}^{1/\rho}) = \sum_{j\in J} \pi_j u'(c_{i,t+1}^j) \approx \sum_{j\in J} \pi_j u'(c_{i,t}R_{i,t,t+1}^{1/\rho}) + \sum_{j\in J} \pi_j (c_{i,t+1}^j - c_{i,t}R_{i,t,t+1}^{1/\rho}) u''(c_{i,t}R_{i,t,t+1}^{1/\rho})$$
$$0 = E_t[c_{i,t+1}] - c_{i,t}R_{i,t,t+1}^{1/\rho}.$$

innovation that is linear in the current transitory and permanent income shocks:

$$\Delta ln(c_{i,t+1}) \approx \underbrace{\Gamma_{t+1}}_{\text{deterministic}} + \underbrace{\phi^{\varepsilon}\varepsilon_{i,t} + \phi^{\eta}\eta_{i,t}}_{\text{uncorrelated with past shocks}}.$$

2 Model and estimator

2.1 Statistical model

The statistical model that I assume encompasses the standard life-cycle model as a special case. In particular, I do not require that a household solves a maximization problem to choose its consumption. The estimating restrictions are as follows.

Log-income growth Log-income is a transitory-permanent income process, the sum of a permanent component that evolves as a random walk and of a transitory component that evolves as an $MA(q)$. The order q of the MA process is established empirically: in the dataset I consider, the covariance between current log-income growth and future log-income growth is no longer statistically significant after $t+2$, so that $q=1$ and I denote $\theta_1 = \theta$. Log-income can also depend linearly on demographic characteristics and other shocks $\zeta^y_{i,t}$ that may capture measurement error. It implies that the log-income growth of household i at period t detrended from the linear effect of its demographic characteristics z, denoted $\Delta ln(\tilde{y}_{i,t})$, is a linear function of the current permanent shock, the current transitory shock, and the past transitory shocks up to period $t-q-1 = t-2$:

$$\Delta ln(\tilde{y}_{i,t}) = \Delta ln(y_{i,t}) - \Delta\kappa_t z_{i,t} = \eta_{i,t} + \varepsilon_{i,t} - (1-\theta)\varepsilon_{i,t-1} - \theta\varepsilon_{i,t-2} + \zeta^y_{i,t} - \zeta^y_{i,t-1}.$$

Log-consumption growth The log-consumption growth of a household i at period t detrended from the linear effect of its demographic characteristics z, denoted $\Delta ln(\tilde{c}_{i,t})$, is a flexible function of the current and past income shocks ε and η it has received:

$$\Delta ln(\tilde{c}_{i,t}) = \Delta ln(c_{i,t}) - \Delta\delta_t z_{i,t} = f(\varepsilon_{i,t}, \eta_{i,t}, \varepsilon_{i,t-1}, \eta_{i,t-1}, ..., \zeta^c_{i,t}, ...).$$

The function $f(.)$ relating income shocks to consumption can be non-linear, and can also depend on consumption-specific shocks ζ^c. As in the case of income, the shocks ζ^c can represent measurement error.

Distributional assumptions I make the following assumptions about the distributions of the variables:

- the demographic characteristics of a household are independent of the current and past shocks it has received ;
- the different shocks received by a household are drawn independently of one another and independently over time;
- either (i) the distribution from which a household draws its transitory shock is normal, and the relation between log-consumption growth and the shocks is unrestricted; (ii) the moments of order higher than two of the distribution are zero, and the third derivative of log-consumption growth with respect to a transitory shock is zero.[14]; (iii) the distribution is unrestricted, but the relation between log-consumption growth and a transitory shock is linear and the effect of a transitory shock is additively separable from the effect of the other current and past shocks (it is the set of assumptions made by BPP who linearize the innovation to log-consumption growth around small income shocks);

Even when transitory shocks are normally distributed, log-income is not necessarily normal, because its distribution also depends on the distributions of demographic characteristics, permanent shocks, and other uncorrelated shocks, about which I remain agnostic. For the standard error to be consistently estimated, I assume that shocks are drawn independently between households, though I allow for arbitrary within-household correlations. There is no assumption imposing uniformity in the distributions from which households draw their shocks, since the estimation procedure is robust to heteroskedasticity: different households can draw their shocks from different distributions, and the same household can draw its shocks from different distributions over time.

Household information The model does not impose that households know about their income process, nor that they can distinguish between the permanent and transitory shocks they receive. If they do not, they will simply respond in the same way to a permanent shock as they do to a transitory shock.

Average elasticity Under these distributional assumptions, I prove in Appendix B that the ratio of the covariance between log-consumption growth and a transitory shock over the variance of this shock coincides with the average elasticity of consumption to

[14]This case only differs from the normality case in that it imposes that the even moments of the distribution be zero instead of being multiples of the variance, and in that it is restrictive regarding the consumption function.

a transitory shock in the population, denoted ϕ^{ε}:

$$\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})}{var(\varepsilon_{i,t})} = E\left[\frac{d\Delta ln(\tilde{c}_{i,t})}{d\varepsilon_{i,t}}\right] = E[\phi^{\varepsilon}_{i,t}].$$

2.2 Identification: instrumenting with future income growth

When the realizations of the shocks ε and η are observed, typically in the context of natural experiments, it is straightforward to measure the covariance $cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})$ and variance $var(\varepsilon_{i,t})$, and to estimate the elasticity to a transitory shock. In survey data, however, the realizations of the shocks ε are not directly accessible. Only total income y is reported, and current log-income growth is driven by the realizations of several different shocks: the current permanent shock, the current transitory shock, and the past transitory shocks.

The solution I use is to instrument the effect of current log-income growth with future log-income growth, which correlates with the realization of the current transitory shock but not with the realization of the current permanent shock. Indeed, a positive transitory income shock raises log-income at t, increasing log-income growth from $t-1$ to t, but as its effect dissipates it does not raise log-income by as much at $t+1$ and no longer raises it at $t+2$, decreasing log-income growth from t to $t+1$ and from $t+1$ to $t+2$. Contrary to this, as a positive permanent shock raises log-income once and for all, so it increases log-income growth from $t-1$ to t but does not decrease it at subsequent periods. In addition, log-income growth at $t+2$ correlates with the realization of the current transitory shock but not with the realizations of any past transitory shocks, whose effects are already fully dissipated. Thus, the covariance between log-consumption growth at t and log-income growth at $t+2$ is exclusively driven by the realization of the transitory shock at t:

$$cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2})) = \theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}).$$

Similarly, the covariance between log-income growth at t and log-income growth at $t+2$ is driven by the realization of the transitory shock at t:

$$cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2})) = \theta var(\varepsilon_{i,t}).$$

An estimator of the elasticity of consumption to a transitory shock that is robust to the presence of a correlation between log-consumption growth and past shocks is:

$$\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon}.$$

2.3 Bias from neglecting the correlation with past shocks

Neglecting the presence of a correlation between log-consumption growth and the realizations of past shocks leads a researcher to use additional instruments that are endogenous when such a correlation is present. Indeed, when past income shocks have no effect on current log-consumption growth, it is possible to use also the opposite of future log-income growth at $t+1$ as an instrument, to get more identifying moments. When past shocks do have an effect, log-income growth at $t+1$ correlates with log-consumption growth through both the current transitory shock and the immediately past transitory shock, so that using it as an additional instrument without acknowledging the effect of the past shock on log-consumption growth yields a bias. In particular, when a past transitory shock correlates negatively with log-consumption growth, because it correlates positively with the opposite of future log-income growth, it reduces the covariance between log-consumption growth and the additional instrument:

$$\begin{cases} cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2})) & = \theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}), \\ cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+1})) & = (1-\theta) cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}) + \theta \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1})}_{<0 \text{ (neglected)}}. \end{cases}$$

The underestimation of this covariance induces an underestimation of the elasticity of consumption to a transitory shock. The two expressions simultaneously identifying the elasticity consumption to a transitory shock are:

$$\begin{cases} \hat{\phi}^{\varepsilon}_{BPP} & = \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon}, \\ \hat{\phi}^{\varepsilon}_{BPP} & = \frac{\theta}{1-\theta} \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(y_{i,t+1}))}{cov_t(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon} + \frac{\theta}{1-\theta} \underbrace{\frac{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1})}{var_t(\varepsilon_{i,t})}}_{<0 \text{ (neglected)}} < \phi^{\varepsilon}. \end{cases}$$

The first identifying expression is unbiased, it is the same as in the robust estimator, but the second expression underestimates the elasticity. The negative effect of the past transitory shock on the covariance between log-consumption growth and the additional instrument is erroneously attributed to the fact that log consumption responds less to the current transitory shock than it really does.

When the process for transitory income is an MA(0), there is only one instrument available, future income growth at $t+1$, and it yields an unbiased estimate because it captures only the dissipation of the transitory shock that realized at t. Erroneously assuming an MA(0) process when the true process is an MA(1), however, leads to an even more severe bias than when the correlation with past shocks is neglected. It induces a researcher to use income growth at $t+1$ as the only instrument:

$$\begin{aligned} & cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+1})) \\ & \qquad = cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}) \underbrace{-\theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})}_{<0} + \underbrace{\theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1})}_{<0}, \end{aligned}$$

and it misses the term $-\theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})$ in addition to missing $\theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1})$. Such an erroneous assumption would also generate a bias in the estimation of the variance of the transitory shock.

3 Results

I implement the robust estimator in data from the Panel Study of Income Dynamics (PSID) between 1978 and 1992, which contains longitudinal information on a representative sample of US households surveyed every year. This PSID data is combined with consumption data imputed from the Consumer Expenditure Survey (CEX) over the same period. I describe the sample selection and the definitions of variables in Appendix C. I implement the estimator in this dataset with a generalized method of moment that is detailed in Appendix D.

3.1 Estimating moments

Table 1: Covariances between $\Delta ln(c)$ or $\Delta ln(y)$ and past, present and future $\Delta ln(y)$

Covariances	$\Delta ln(\tilde{y}_{i,t-1})$	$\Delta ln(\tilde{y}_{i,t})$	$\Delta ln(\tilde{y}_{i,t+1})$	$\Delta ln(\tilde{y}_{i,t+2})$	$\Delta ln(\tilde{y}_{i,t+3})$
$cov(\Delta ln(\tilde{c}_{i,t}), .)$	-0.0029	0.0128	-0.0004	-0.0031	0.0010
	(0.0017)	(0.0019)	(0.0018)	(0.0015)	(0.0015)
$cov(\Delta ln(\tilde{y}_{i,t}), .)$	-0.0252	0.0785	-0.0258	-0.0058	0.0006
	(0.0018)	(0.0032)	(0.0016)	(0.0013)	(0.0014)
Obs.	7,578	8,958	8,958	8,958	7,539

Note: Standard errors in parentheses are adjusted for arbitrary within-household correlations and heteroskedasticity. The first two lines report the covariances between log-consumption growth or log-income growth and past, present or future log-income growth.

Covariances of log-consumption growth Before looking at the elasticity estimates, Table 1 presents the value of the moments used for estimation. The first line shows the covariances between log-consumption growth and past, current, and future log-income growth. They are consistent with my statistical model, but inconsistent with a model that does not allow for a correlation between log-consumption growth and past income shocks. First, the covariance between log-consumption growth and future log-income growth at $t+1$ is smaller in absolute value than its covariance with log-income growth at $t+2$: it is -0.0004 at $t+1$, not statistically significant, but -0.0031 at $t+2$, statistically significant at 5%. A correlation between log-consumption growth and past shocks can account for this stylized fact: log-consumption growth covaries negatively with log-income growth at $t+1$ through the current transitory shock but can also covary positively with it through the past transitory shock, reducing the magnitude of the covariance; log-consumption growth covaries with log-income growth at $t+2$ only through the current transitory shock, generating a strongly negative covariance. On the contrary, a version of this model that does not allow for a correlation with past shocks predicts that the covariance between log-consumption growth and future log-income growth at $t+1$ must be larger in absolute value than its covariance with log-income growth at $t+2$ by a factor of $\frac{1-\theta}{\theta}$. Second, Table 1 shows that the covariance between log-consumption growth and past log-income growth is large, -0.0029, and statistically significant at 10%. This is conceivable in a model that allows for a correlation with past shocks, since past log-income growth is a sum of past shocks, but is incompatible with a model that does not allow log-consumption growth to correlate with past shocks.[15]

Covariances of log-income growth The autocovariances of log-income growth presented in the second line of Table 1 suggest that the transitory component of income is an MA(1). If the transitory component was an MA(0), the covariance between log-income growth at t and log-income growth at $t+2$ should not be statistically different from zero, while it is. If the transitory component was an MA(2), the covariance between log-income growth at t and log-income growth at $t+3$ should be statistically different from zero, while it is not. Also, if permanent income was not a random walk but an AR(1) with a coefficient different from one, the autocovariances between log-income growth at t and at all future periods should be statistically different from zero, while they stop being significant after two periods. I still examine the case in which permanent income is an AR(1) with a coefficient just slightly below one.

[15] Note that, even in the presence of a non-zero correlation with past shocks, this covariance could be close to zero or imprecisely estimated because past log-income growth is $\Delta ln(y_{i,t-1}) = \eta_{t-1} + \varepsilon_{t-1} - \theta\varepsilon_{t-2} - (1-\theta)\varepsilon_{t-3}$, so the covariances with the different shocks could cancel off.

3.2 Elasticity of consumption to a transitory shock

Table 2: Elasticity ϕ^{ε} and MPC lower bound $\underline{MPC}^{\varepsilon}$

	Robust	Non-robust	BPP (non-robust)
ϕ^{ε}	0.539	0.052	0.046
	(0.274)	(0.033)	(0.040)
$\underline{MPC}^{\varepsilon}$	0.450	0.043	0.035
	(0.229)	(0.027)	(.030)
$\underline{MPC}^{\varepsilon\ total}$ with $\theta = 0.3$	0.343	0.033	0.028
	(0.175)	(0.021)	(0.024)
$\underline{MPC}^{\varepsilon\ total}$ with $\theta = 0.9$	0.237	0.023	0.019
	(0.120)	(0.014)	(0.016)
Obs.	8,958	8,958	12,041
Moments used	(1)	(1), (2)	$(1)_t$, $(2)_t$ $\forall t$ + others

Note: Standard errors in parentheses are adjusted for arbitrary within-household correlations and heteroskedasticity. The first line reports GMM estimates of the average elasticity of nondurable consumption to a transitory shock on net income. The estimating moments are:

$$0 = cov(\Delta ln(\tilde{c}_{i,t}), \Delta ln(\tilde{y}_{i,t+2})) - \phi^{\varepsilon} cov(\Delta ln(\tilde{y}_{i,t}), \Delta ln(\tilde{y}_{i,t+2})) \quad (1)$$

$$\underbrace{0 = cov(\Delta ln(\tilde{c}_{i,t}), \Delta ln(\tilde{y}_{i,t+1})) - ((1-\theta)/\theta)\phi^{\varepsilon} cov(\Delta ln(y_{i,t}), \Delta ln(\tilde{y}_{i,t+2})))}_{\text{non-robust to a correlation with past shocks}} \quad (2)$$

The first column presents an estimate based on moment (1); the second column an estimate based on moments (1) and (2), where the value of θ is externally imposed as the value used in BPP of $\theta = 0.1132$; the third column presents the results of the original BPP estimator, based on moments (1) and (2) taken conditionally on the period t and on additional moments that estimate other parameters. The second line reports a lower bound for the MPC out of the change in current income caused by a transitory shock. The third and fourth lines report lower bounds for the MPC out of total the net present value change in income caused by a transitory shock for $\theta = 0.3$ and $\theta = 0.9$. The fourth line reports the number of household-years for which at least one estimating moment in which ϕ^{ε} appears is observed.

Robust estimator The first column of Table 2 reports the results obtained with my robust estimator, which is based on a moment that holds even in the presence of a correlation between log-consumption growth and past income shocks. Using this estimator, the average elasticity of nondurable consumption to a transitory shock on net income is large, with a point estimate of 0.54, statistically significant at 5%. It means that a transitory shock that raises current income by 10% and next period income by $\theta \times 10\%$ leads to a 5.4% increase in current nondurable consumption over the following year, on average in the sample.

Non-robust estimators The next two columns show the results obtained with estimators that are not robust to the presence of a correlation between log-consumption growth and past income shocks, because they rely on moment (2), an expression that neglects a term that is non-zero in the presence of such a correlation. These two estimators yield much smaller estimates of the average elasticity of consumption to a transitory income shock than the robust estimator. The second column features the results from a simple non-robust estimator that is identical to the robust estimator except that it additionally relies on moment (2). The associated estimate of the elasticity of consumption to a transitory shock is 0.05, much below the estimate of 0.54 obtained with the robust estimator, and not statistically significant. The fact that the point estimate is substantially smaller is consistent with my theoretical prediction that using moment (2) for estimation induces a downward bias in the measure of the elasticity of consumption to a transitory income shock when log-consumption growth and past transitory shocks correlate negatively. The third column presents the results obtained with the original BPP estimator, which uses moments (1) and (2) as well, but differs from the robust estimator and the simple non-robust estimator on other grounds because it relies on additional moments and considers all moments conditionally on the period. The point estimate of the average elasticity of consumption to a transitory income shock is similar to that of the simple non-robust estimator, at 0.05, and not statistically significant.[16] This suggests that, although additional estimating moments are used in the BPP estimator, they serve to identify other parameters and the elasticity to a transitory shock remains identified mainly from moments (1) and (2).[17]

Marginal Propensity to Consume (MPC^{ε}) Semi-structural estimators measure the elasticity of consumption to a transitory shock $\phi^{\varepsilon}_{i,t} = \frac{1}{c_{i,t}} \frac{dc_{i,t}}{d\varepsilon_{i,t}}$, that is, the percentage change in current consumption associated with a one unit transitory income shock,

[16] To make things comparable, the results presented in the third column are obtained by implementing the original BPP estimator in the same measures of log-consumption growth and log-income growth as used in the first two columns, which differ slightly from the original measures used by BPP because of my additionally interacting demographic variables by cohort. Implementing the original BPP estimator in the original measures of log-consumption growth and log-income growth gives the very similar estimate of 0.053.

[17] For the first two estimators, the observations are the household-years such that $\Delta ln(\tilde{c}_{i,t})$, $\Delta ln(\tilde{y}_{i,t})$, and $\Delta ln(\tilde{y}_{i,t+2})$ are simultaneously observed. For the BPP estimator, the observations are the household-years such that $\Delta ln(\tilde{c}_{i,t})$ and $\Delta ln(\tilde{y}_{i,t})$ are simultaneously observed, since the estimator uses the household-years such that $\Delta ln(\tilde{c}_{i,t})$, $\Delta ln(\tilde{y}_{i,t})$, and $\Delta ln(\tilde{y}_{i,t+1})$ or $\Delta ln(\tilde{y}_{i,t+2})$ are observed to estimate the covariance between log-consumption growth and a transitory shock, and it uses the household-years such that $\Delta ln(\tilde{c}_{i,t})$ and $\Delta ln(\tilde{y}_{i,t})$ are observed to estimate the variance of the transitory shock.

which causes a 100% change in current income and a $\theta \times 100\%$ change in future income.[18] On the contrary, natural experiments typically measure the MPC out of current income, that is, the level change in current consumption caused by a level change in current income: $MPC_{i,t}^{\varepsilon} = \frac{dc_{i,t}}{dy_{i,t}} = \frac{d\varepsilon_{i,t}}{dy_{i,t}} \frac{dc_{i,t}}{d\varepsilon_{i,t}} = \frac{1}{y_{i,t}} \frac{dc_{i,t}}{d\varepsilon_{i,t}}$. The relation between the elasticity of consumption and the MPC out of the change in current income caused by a current transitory shock is:

$$MPC_{i,t}^{\varepsilon} = \frac{c_{i,t}}{y_{i,t}} \phi_{i,t}^{\varepsilon}.$$

A difficulty is that I do not measure the individual elasticities $\phi_{i,t}^{\varepsilon}$ but only the average elasticity in the sample $\phi^{\varepsilon} = E[\phi_{i,t}^{\varepsilon}]$. Under the assumption that the individual elasticities are all equal (and thus equal to the average elasticity), the average MPC is the product of the average elasticity and the average ratio of consumption over income. Under the alternative assumption that individual elasticities are not necessarily all equal, but are such that the households with the highest elasticities are on average those with the highest ratios of consumption over income, this product is a lower bound for the average MPC:

$$\underline{MPC^{\varepsilon}} = E[\frac{c_{i,t}}{y_{i,t}}]\phi^{\varepsilon} \leq E[\frac{c_{i,t}}{y_{i,t}} \phi_{i,t}^{\varepsilon}] = MPC^{\varepsilon}.$$

Because the average ratio of consumption over income is smaller than one, the lower bound for the average MPC is smaller than the average elasticity. I find that a household consumes on average at least 45% of the change in its current income caused by a transitory shock. The lower bounds for the average MPC measured with non-robust estimators are small, at 4%, and not statistically significant.

Marginal Propensity to Consume out of Total Income Change ($MPC^{\varepsilon\, total}$) While the transitory shocks identified in survey data are found to evolve as MA(1), affecting both current and future income, most of the shocks considered in natural experiments are purely transitory, affecting current income only. Considering the limit case in which a change in future income has no effect on current consumption, for instance because of borrowing constraints or myopic behavior, this difference in persistence is not a problem, and the MPC out of the change in current income caused by an MA(1)

[18]The change in current income caused by a transitory shock is $\frac{dy_{i,t}}{d\varepsilon_{i,t}} = \frac{de^{p_{i,t}} e^{\varepsilon_{i,t}} e^{\theta \varepsilon_{i,t-1}}}{d\varepsilon_{i,t}} = y_{i,t}$: a one unit increase in $\varepsilon_{i,t}$ raises current income by 100%, that is, doubles it. The change in future income caused by a transitory shock is $\frac{dy_{i,t+1}}{d\varepsilon_{i,t}} = \frac{de^{p_{i,t+1}} e^{\varepsilon_{i,t+1}} e^{\theta \varepsilon_{i,t}}}{d\varepsilon_{i,t}} = \theta y_{i,t+1}$: a one unit increase in $\varepsilon_{i,t}$ raises current income by $\theta \times 100\%$.

transitory shock also measures how households would respond to a purely transitory shock. Considering the opposite case in which a change in future income has exactly the same effect on current consumption as a change in current income, it is the MPC out of the total net present value change in income caused by an MA(1) transitory shock that measures how households would respond to a purely transitory shock: $MPC^{\varepsilon\, total}_{i,t} = \frac{dc_{i,t}}{dy_{i,t}} \frac{1}{y_{i,t}+\theta/(1+r)y_{i,t+1}}$.[19] The relation with the elasticity is:

$$MPC^{\varepsilon\, total}_{i,t} = \frac{c_{i,t}}{y_{i,t} + (\theta/(1+r))y_{i,t+1}} \phi^{\varepsilon}_{i,t}.$$

Measuring this MPC requires taking a stand on the value of the interest factor $(1+r)$ and on the persistence of the transitory shock θ. I select $(1+r) = 1.03$ and I present results for two limit values of θ, starting at $\theta = 0.3$ because Meghir and Pistaferri (2004) find that a lower bound for the coefficient of the MA process is $\theta = 0.26$.[20] A lower bound for the average MPC out the total net present value change in income caused by a transitory shock is:

$$MPC^{\varepsilon\, total} = E[\frac{c_{i,t}}{y_{i,t} + \theta/(1+r)y_{i,t+1}} \phi^{\varepsilon}_{i,t}] \geq E[\frac{c_{i,t}}{y_{i,t} + \theta/(1+r)y_{i,t+1}}]\phi^{\varepsilon} = \underline{MPC}^{\varepsilon\, total}.$$

The $MPC^{\varepsilon\, total}$ decreases with the value of θ because a higher θ implies that the total net present value change in income caused by a transitory shock is larger for the same change in current consumption. Assuming that a current transitory shock only multiplies future income by 0.3, a lower bound for $MPC^{\varepsilon\, total}$ is 0.34: a household consumes at least 34% of the total net present value gain in income caused by a transitory shock over the following year. Assuming that a transitory shock is almost fully passed on to the next period ($\theta = 0.9$), a lower bound for $MPC^{\varepsilon\, total}$ is 0.24, so a household consumes at least 24% of the total net present value gain in income caused by a transitory shock over the following year.

Comparison with the literature on natural experiments How do this elasticity and these MPCs compare with the results derived from natural experiments? Studying in-

[19] I compute the MPC out of the realized change in current and future income caused by a current transitory shock, and not the MPC out of the expected change $y_{i,t} + (\theta/(1+r))E[y_{i,t+1}]$, because I do not observe expected income $E[y_{i,t+1}] = e^{p_t}E_t[e^{\eta_{i,t+1}}]E_t[e^{\varepsilon_{i,t+1}}]$.

[20] The value of $\theta = 0.1132$ that I plug in the simple non-robust estimator is the value that is used in the BPP estimator, so it is under this value that the bias they make must be calculated. However, it is not necessarily the true value of θ: although BPP obtain it by estimating the income process, they do not allow for measurement error in their estimation while Meghir and Pistaferri (2004) show that the variance of measurement error confounds the value of θ. They re-estimate θ by relying on an external estimate of measurement error.

creases in take-home pay, Parker (1999) finds that the average elasticity of nondurable consumption over the next three months out of a temporary increase in take-home pay caused by a change in social security taxes is 0.54, significant at 1%. Souleles (2002) estimates the MPC of nondurable consumption over the next three months out of a change in take-home pay induced by the Reagan tax cuts to be 0.66, significant at 5%. Baker and Yannelis (2017) and Gelman, Kariv, Shapiro, Silverman, and Tadelis (2018) estimate the MPC of nondurable consumption and of total credit card spending over the next two weeks out of the temporary decrease in take-home pay caused by the 2013 government shutdown to be 0.39 and 0.58, both statistically significant at 1%. Studies of tax refunds, tax rebates, and stimulus program receipts, that are mailed directly to the households, obtain MPCs of nondurable expenditures over the next three months that are between 0.09 and 0.37.[21] Kan, Peng, and Wang (2017) consider the effect of a shopping voucher program, and find that it stimulates reported total expenses above what would have been spent otherwise by 0.24 of the voucher value over the next three months, significant at 1%. Fagereng, Holm, and Natvik (2018) find that the MPC of total spending (measured in a broad way as the difference between income growth and wealth growth) out of a small lottery prize (below $\$2,150$) over the next year is 1.01, significant at 1%. Thus, two characteristics seem robust to the idiosyncrasies of the shocks considered in these different natural experiments: (i) the MPC of nondurable consumption out of an unexpected transitory shock is statistically significant (ii) its value over the year following the shock is at least 0.10. The MPC derived from a robust estimator is consistent with both stylized facts, while the MPCs derived from non-robust estimators conflict with both.

Robustness I cannot directly estimate the moments of the distribution of the transitory shocks in the sample, but I report moments that are proportional to them by a factor θ, shown in section D of the online appendix. They are consistent with any of the three alternative sets of assumptions that I require: none of the odd moments are statistically different from zero, and none of the moments of order higher than two are statistically different from zero, except for the fourth moment that is statistically significant at 10%.

I check that the findings are robust to the choice of demographic characteristics, interactions, clusters, and measures of consumption and income that I make. Results are reported in section D of the online appendix. Incidentally, under the assumptions of the statistical model, the rationale for detrending log-income and log-consumption growth from the effect of demographic characteristics is only to avoid serial correlation in de-

[21] See Souleles (1999), Johnson, Parker, and Souleles (2006), Parker, Souleles, Johnson, and McClelland (2013), Misra and Surico (2014), and Agarwal and Qian (2014).

mographic changes, and possibly to increase precision by eliminating noise.[22] Consistent with that, I find that excluding groups of demographic characteristics from the set of detrending variables has little effect on the point estimate obtained. Even in the limit case in which the variables are only detrended by year dummies, the elasticity is 0.44. The results for food expenditures, total expenditures (nondurables and durables), and total expenditures plus expenditures on education and health (not considered durables as they could capture investment) are consistent with comparable natural experiments findings. The response of consumption to shocks on measures of income that exclude taxes, or exclude both taxes and transfers, is smaller, in line with the role that taxes and transfers are expected to play. It means that, when households receive a shock on their income before taxes and transfers, they anticipate adjustments in taxes and transfers that will reduce the magnitude of the shock, so they respond less to it than they do to a shock on income after taxes and transfers.

Finally, I compute the average elasticity of future consumption to a current transitory shock, to get a sense of the dynamics of the response. The marginal propensity to consume over the two years following a transitory shock is statistically significant at 5% for any $\theta \geq 0.4$. For this range of θ, its point estimate is between 0.24 and 0.39, as reported in section D of the online appendix.

[22]Formally, by assumption, the effect of demographic characteristics on log-consumption growth and log-income growth is linear and additively separable from the effect of the income shocks and other shocks: $\Delta ln(c_{i,t}) = \Delta ln(\tilde{c}_{i,t}) + \Delta\delta_t z_{i,t}$ and $\Delta ln(y_{i,t}) = \Delta ln(\tilde{y}_{i,t}) + \Delta\kappa_t z_{i,t}$; also, by assumption, the demographic characteristics are independent of the realizations of the current and past shocks; thus, the only difference between using detrended and non-detrended variables is the presence of the terms $cov(\Delta\delta_t z_{i,t}, \Delta\kappa_{t+2} z_{i,t+2})$ and $cov(\Delta\kappa_t z_{i,t}, \Delta\kappa_{t+2} z_{i,t+2})$, which capture serial correlation.

3.3 Heterogeneity across household characteristics

Table 3: Elasticity ϕ^{ε} and MPC lower bounds $\underline{MPC}^{\varepsilon}$ across different subgroups

	Financial income ($\leq/>$ 300\$ per year)		Wage rate ($\leq/>$ 13\$ per hour)		Age ($\leq/>$ 44 years old)	
	Low	High	Low	High	Low	High
ϕ^{ε}	0.608	0.493	0.601	0.302	0.849	0.326
	(0.492)	(0.299)	(0.300)	(0.603)	(0.503)	(0.313)
$\underline{MPC}^{\varepsilon}$	0.525	0.396	0.626	0.160	0.696	0.278
	(0.425)	(0.240)	(0.312)	(0.320)	(0.412)	(0.267)
Obs.	4,556	4,402	5,323	3,635	5,030	3,928

Note: Standard errors in parentheses are adjusted for arbitrary within-household correlations and heteroskedasticity. The first line reports GMM estimates of the average elasticity of nondurable consumption to a transitory shock on net income, computed within different subgroups of household-years. The second line reports a lower bound for the MPC out of the change in current income caused by a transitory shock. The third line reports the number of household-years for which the estimating moment is observed.

Table 3 presents the estimates obtained when running the robust estimator separately on different subgroups of the sample.[23] I emphasize two facts. First, the point estimates are all above about 0.30, and the lower bounds for the MPC all above 0.15. This suggests that all subgroups of the sample respond quite substantially to a transitory income shock, although the estimates are not all statistically significant, possibly because of the reduced sample sizes. Second, the point estimates are larger for households with lower financial income, a lower wage rate of the head, and a younger head. The average elasticities are all above 0.60 for these three subgroups. The average elasticity of the households with a young head is particularly high. The lower bounds for the MPC are all above 0.50 for these three subgroups. Possible interpretations for this result are that these households are less likely to own liquid assets, so they are less able to smooth liquidity shocks, or that they are facing more uncertainty (young households in particular), two features that can raise the elasticity of consumption in life-cycle models.[24] As there are no direct questions about liquid assets in the 1978-1992 PSID, it is difficult to test their effect on the response to a transitory shock.

For the three breakdowns considered, the weighted sum of the two subgroup elasticities coincides quite closely with the elasticity over the full sample, suggesting that

[23]The variables are detrended over the whole sample and not over each subsample, to keep the effect of demographic characteristics constant and make the comparison between them more direct.

[24]See Deaton (1991), Kimball (1990), Carroll and Kimball (1996), Kaplan and Violante (2014).

the elasticity estimate over the whole sample is the average elasticity in the sample, as would be the case if my distributional assumptions hold.

3.4 Depreciation of the permanent shock

I consider a more general income process in which permanent income is not necessarily a random walk, but simply an AR(1) with coefficient ρ:

$$p_t = \rho p_{t-1} + \eta_t.$$

This means that a permanent shock η_t still affects the value of permanent income at each period in the future until the rest of the household's lifetime, but the effect of this permanent shock now depreciates at a rate $(1-\rho)$ instead of affecting all values of future permanent income in a the same way. As a result, the log-income growth of a household at t depends, not only on the past transitory shocks it has received at $t-1$ and $t-2$, but also on all the permanent shocks it has received in the past:

$$\begin{aligned}\Delta ln(\tilde{y}_t) =& \eta_t - (1-\rho)\eta_{t-1} - (1-\rho)\rho\eta_{t-2} - ... - (1-\rho)\rho^{t-2}\eta_1 - (1-\rho)\rho^{t-1}p_0 \\ &+ \varepsilon_t - (1-\theta)\varepsilon_{t-1} - \theta\varepsilon_{t-2}.\end{aligned}$$

The instrument, future log-income growth at $t+2$, no longer identifies only the effect of the current transitory shock because it covaries with current log-consumption growth through the current transitory shock but also through the current and past permanent shocks. Since the correlation between log-consumption growth and past permanent shocks is undetermined, it is difficult to predict the direction of the bias caused by this depreciation. Yet, Kaplan and Violante (2010) show that if the value of ρ is known, it is possible to obtain a consistent estimator by substituting log-income growth with its quasi-difference $\Delta^\rho ln(\tilde{y}_{i,t}) = ln(\tilde{y}_t) - \rho ln(\tilde{y}_{t-1})$ in the estimating moments. The estimator is:

$$\hat{\phi}^{\varepsilon,\rho} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), \Delta^\rho ln(\tilde{y}_{i,t+2}))}{cov(\Delta^\rho ln(\tilde{y}_{i,t}), \Delta^\rho ln(\tilde{y}_{i,t+2}))} = \phi^\varepsilon.$$

Table 4: Elasticity $\phi^{\varepsilon,\rho}$ and MPC lower bound $\underline{MPC}^{\varepsilon,\rho}$ when permanent shocks depreciate

Depreciation	$\rho = 1$	$\rho = 0.97$	$\rho = 0.94$	$\rho = 0.91$
$\phi^{\varepsilon,\rho}$	0.539	0.549	0.580	0.642
	(0.274)	(0.289)	(0.318)	(0.371)
$\underline{MPC}^{\varepsilon,\rho}$	0.450	0.458	0.483	0.535
	(0.229)	(0.241)	(0.265)	(0.309)
Obs.	8,958	8,958	8,958	8,958

Note: The first line reports GMM estimates of the average elasticity of non-durable consumption to a transitory shock on net income, computed under the assumption that permanent shocks depreciate at a rate $(1-\rho)$. The second line reports a lower bound for the MPC out of the change in current income caused by a transitory shock. The third line reports the number of household-years for which the estimating moment is observed.

Table 4 presents the results obtained with such an estimator, for different values of ρ. The first column corresponds to the baseline case in which $\rho = 1$. The point estimate of the elasticity of consumption then increases as ρ decreases below one: it moves from 0.54 to 0.55, 0.58, and 0.64 when ρ moves from 1 to 0.97, 0.94, and 0.91. It means that, if in fact permanent shocks depreciate over time, the robust estimator is conservative and underestimates the elasticity.

3.5 Anticipation of the shocks

I generalize the statistical model by allowing part of the realizations of the permanent and transitory shocks at t to be anticipated at previous periods $t-s$ and $t-k$:

$$\eta_{i,t} = \eta_{i,t}^{surp} + \eta_{i,t-s}^{ant,t},$$
$$\varepsilon_{i,t} = \varepsilon_{i,t}^{surp} + \varepsilon_{i,t-k}^{ant,t}.$$

Each type of shock writes as the sum of a surprise component and an anticipated component whose value realizes at t but is known at $t-s$ or $t-k$. The anticipated component is uncorrelated with the current and past surprise components and with the past anticipated components.

Earliest anticipation period From the moments presented in Table 1, log-consumption growth at t no longer covaries with future log-income growth at $t+3$ or later, which is informative about how early each type of shock can be anticipated. Indeed, in the presence of anticipation, the covariance between log-consumption growth and future log-income growth at $t+3$ is:

$$
\begin{aligned}
0 = cov(\Delta ln(\tilde{c}_{i,t}), \Delta ln(\tilde{y}_{i,t+3})) = & \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \eta^{ant,t+3}_{i,t+3-s})}_{\neq 0 \text{ if } s>2} + \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon^{ant,t+3}_{i,t+3-k})}_{\neq 0 \text{ if } k>2} \\
& - (1-\theta) \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon^{ant,t+2}_{i,t+2-k})}_{\neq 0 \text{ if } k>1} - \theta \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon^{ant,t+1}_{i,t+1-k})}_{\neq 0 \text{ if } k>0}.
\end{aligned}
$$

Apart from a knife-edge case in which the effects would perfectly compensate each other, the earliest period at which a permanent shock can be anticipated that is consistent with the covariance above being zero is $s=2$. The earliest period at which a transitory shock can be anticipated is $k=0$, that is, a transitory shock cannot be anticipated beforehand.[25]

Bias from anticipation Applying the robust estimator to a model in which permanent shocks are partly anticipated gives:

$$
\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon} - \underbrace{\frac{cov(\Delta ln(\tilde{c}_{i,t}), \eta^{ant,t+2}_{i,t+2-s})}{\theta var(\varepsilon_{i,t})}}_{=0 \text{ if } s=1 / >0 \text{ if } s=2} \leq \phi^{\varepsilon}.
$$

The denominator of the estimator is unaffected because whether a household can anticipate future income shocks or not does not affect its current log-income growth, which only correlates with future log-income growth through the realization of the transitory shock. The numerator can be affected, however, because information about its future income can impact the current consumption decision of a household, inducing a correlation between current log-consumption growth and future log-income growth through the anticipated component of future log-consumption growth. If the anticipated component is known only one period in advance, though, the numerator remains unaffected: no component of the instrument is anticipated yet at t, so the current transitory shock is still the only variable through which current log-consumption growth and the instrument covary. If the anticipated component is known two periods in advance, and if log-consumption is positively correlated with current information about the realization

[25]If households are myopic or constrained, it is possible that they anticipate shocks earlier but simply do not respond to them. In that case, the fact that the covariance is zero does not imply that shocks are not anticipated earlier than $s=2$ and $k=0$. Yet, in that case, an early anticipation of the shocks does not bias the estimation since consumption does not respond to the anticipated component of a shock before it actually realizes.

of future permanent income, then the robust estimator is a conservative measure of the elasticity and underestimates it.

3.6 Serial correlation in measurement error

I consider a more general model in which measurement error can be correlated over time. Measurement error ζ^y is no longer orthogonal to its past values but such that: $\zeta^y_{i,t} = \dot{\zeta}^y_{i,t} + \nu\dot{\zeta}^y_{i,t-1}$, with ν a parameter measuring the strength of the serial correlation. Thus, log-income growth depends on past measurement error $\dot{\zeta}^y$ up to two periods ago:

$$\Delta ln(\tilde{y}_{i,t}) = \eta_{i,t} + \varepsilon_{i,t} - (1-\theta)\varepsilon_{i,t-1} - \theta\varepsilon_{i,t-2} + \dot{\zeta}^y_{i,t} - (1-\nu)\dot{\zeta}^y_{i,t-1} - \nu\dot{\zeta}^y_{i,t-2}$$

The robust estimator is:

$$\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon} \times \underbrace{\frac{\theta var(\varepsilon_{i,t})}{\theta var(\varepsilon_{i,t}) + \nu var(\dot{\zeta}^y_{i,t})}}_{<1} < \phi^{\varepsilon}.$$

The presence of serial correlation in measurement error leads to an overestimation of the variance of the current transitory shock, because the variance of measurement error is mistaken as variance of the transitory shock: the denominator measures $\theta(var(\varepsilon_{i,t}) + (\nu/\theta)var(\dot{\zeta}_{i,t}))$ instead of $\theta var(\varepsilon_{i,t})$. The estimation of the covariance between log-consumption growth and a current transitory shock is unaffected because consumption does not respond to measurement error. Thus, in the presence of serial correlation, the robust estimator is conservative and underestimates the elasticity of consumption to a transitory shock because it overstates the importance of the shocks. If I additionally allow measurement error in consumption to correlate with measurement error in income, the robust estimator would overestimate both the covariance between log-consumption growth and a transitory shock and the variance of the transitory shock, still underestimating the elasticity to a transitory shock as long as the elasticity of consumption to measurement error in income is smaller than its elasticity to a transitory shock and that ν is small when compared to θ.[26]

[26]In that case, the robust estimator measures $\hat{\phi}^{\varepsilon} = \frac{\theta cov(\Delta ln(c_{i,t}),\varepsilon_{i,t}) + \nu cov(\Delta ln(c_{i,t}),\dot{\zeta}^y_{i,t})}{\theta var(\varepsilon_{i,t}) + \nu var(\dot{\zeta}^y_{i,t})} < \frac{cov(\Delta ln(c_{i,t}),\varepsilon_{i,t})}{var(\varepsilon_{i,t})} = \phi^{\varepsilon}$ if $\frac{cov(\Delta ln(c_{i,t}),\dot{\zeta}^y_{i,t})}{var(\dot{\zeta}^y_{i,t})} < \frac{1-\theta}{\nu}\frac{cov(\Delta ln(c_{i,t}),\varepsilon_{i,t})}{var(\varepsilon_{i,t})}$.

4 Related estimation methods

4.1 Estimation of the elasticity to a permanent shock

Method - joint estimation When past shocks have no effect on log-consumption growth, it is possible to use log-income growth between $t-1$ and $t+2$ as an instrument to identify the elasticity of consumption to a permanent shock. Indeed, a transitory shock at t raises log-income at t, at $t+1$ and then dissipates, so it has no impact on log-income growth between $t-1$ and $t+2$. On the contrary, a permanent shock at t raises log-income permanently from period t on, so it does raise log-income growth between $t-1$ and $t+2$.

Bias - joint estimation Yet, when past shocks affect log-consumption growth, this instrument is not exogenous as it does not identify the effect of the current permanent shock separately from the effect of past shocks. In particular, when past transitory shocks correlate negatively with log-consumption growth, they raise the covariance between log-consumption growth and log-income growth between $t-1$ and $t+2$. The overestimation of this covariance translates into an overestimation of the elasticity of consumption to a permanent shock:

$$\hat{\phi}^{\eta}_{BPP} = \phi^{\eta} + \underbrace{\frac{cov(\Delta ln(\tilde{c}_{i,t}), -\varepsilon_{i,t-1} - \theta\varepsilon_{i,t-2})}{var(\eta_{i,t})}}_{>0 \text{ (neglected)}} > \phi^{\eta}.$$

This expression overestimates the elasticity because the increase in the covariance between log-consumption growth and the instrument caused by the past transitory shocks is erroneously attributed to the fact that log-consumption growth covaries more with the permanent shock than it really does.

Method - separate estimation Kaplan and Violante (2010) develop an estimator of the elasticity to a permanent shock that is slightly different. The instrument they use is log-income growth between $t-3$ and $t+2$ (when transitory income is an MA(1) process) or between $t-2$ and $t+1$ (when transitory income is an MA(0), which is their assumption in the simulations), which also covaries with the current permanent shock but not with the transitory shock since a transitory shock has no impact on log-income growth over a window period around the shock.

Bias - separate estimation When past income shocks affect log-consumption growth, this instrument is not exogenous because it covaries with log-consumption growth through the current permanent shock, the past permanent shock at $t-1$, and the past transitory shock at $t-2$. The estimator is:

$$\hat{\phi}^{\eta}_{KV} = \phi^{\eta} + \underbrace{\frac{cov(\Delta ln(\tilde{c}_{i,t}), \eta_{i,t-1})}{var(\eta_{i,t})}}_{\neq 0 \text{ (neglected)}} + \underbrace{\frac{cov(\Delta ln(\tilde{c}_{i,t}), -\varepsilon_{i,t-2})}{var(\eta_{i,t})}}_{> 0 \text{ (neglected)}}.$$

The direction of the bias depends on whether a past permanent shock covaries positively or negatively with current log-consumption growth, which I have not established analytically in the standard life-cycle model: although having received a good permanent shock in the past raises the level of assets and the expected future income, it also makes future income more volatile, so its effect on precautionary saving is undetermined. As a result, this estimator overestimates the elasticity if the effect of a past permanent shock on subsequent log-consumption growth is positive, or negative but small enough to be dominated by the effect of the past transitory shock; it underestimates the elasticity if the effect of a past permanent shock is negative and large enough to dominate the positive effect of the past transitory shock. Running their version of the BPP estimator on simulated data, Kaplan and Violante (2010) find that their estimator overestimates the elasticity of consumption to a permanent shock (slightly when the borrowing constraint is the natural one, greatly when there is no borrowing allowed). This indicates that, for their particular calibration at least, past permanent shocks correlate either positively or negatively but slightly with subsequent log-consumption growth.[27]

4.2 Estimation of Frisch elasticities

Method Building on the BPP estimator, some studies take a step back from net income, considering shocks that can partly be endogenously insured and estimating the Frisch elasticity of consumption to these shocks, holding the marginal utility of wealth constant, to measure the importance of the changes in margins other than wealth caused by the shock. One of the first to do so is Blundell, Pistaferri, and Saporta-Eksten (2016), who consider the effect of a permanent wage shock on consumption, and the extent to which it is insured through adjustments in the labor supply of the household members. The authors estimate the Frisch elasticities by matching empirical estimates of the elasticities of consumption and labor supply to a permanent wage shock, measured with the BPP method, with their theoretical expressions, which depend on the Frisch elasticities,

[27]Because Kaplan and Violante (2010) estimate separately the elasticities to a transitory and to a permanent shock, and not jointly, the bias affecting the measure of the elasticity to a transitory shock cannot be contaminating their estimate, as can be the case in the original BPP estimator.

on the elasticities to shocks on net income (after all insurance has taken place), and on observable institutional parameters (response of taxes in particular). For instance, the theoretical elasticity of consumption to a permanent wage shock is:

$$\phi^{\eta\ BPS}_{\text{wage}} = f(\phi^{\eta\ BPS}_{\text{net. inc.}}, F_{i,t}, I_{i,t}),$$

with $\phi^{\eta\ BPS}_{\text{net. inc.}} = \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{((1+r)^s}/((1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s})$ an approximation of the theoretical elasticity of consumption to a permanent shock on net income, $F_{i,t}$ the vector of Frisch elasticities, $I_{i,t}$ the vector of institutional parameters, and $f(.)$ a functional form derived from the approximated solution of the elasticity to a permanent wage shock in a life-cycle model. The values of $\phi^{\eta\ BPS}_{\text{net. inc.}}$ and $I_{i,t}$ are externally measured and plugged in to estimate $F_{i,t}$. Incidentally, the theoretical elasticity of consumption to a transitory shock on net income is neglected because it is approximated as the share of current income in lifetime expected resources, $\phi^{\varepsilon\ BPS}_{\text{net. inc.}} = y_t/((1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s})$, which is small. The method is more structural than the BPP estimator, since the authors use the theoretical expressions of the elasticities for estimation, and not just the restriction they imply that log-consumption growth is uncorrelated with past shocks.[28]

Bias There are two issues with this method. First, the empirical counterpart of $\phi^{\eta}_{\text{wage}}$ is measured with the BPP estimator of the elasticity to a permanent shock. This estimator is subject to caution as it is not robust to the presence of a correlation with past income shocks, although in the case of permanent shocks I do not quantify the extent of the bias that this correlation causes.

Second, the theoretical expressions $\phi^{\eta\ BPS}_{\text{wage}}$, $\phi^{\eta\ BPS}_{\text{net. inc.}}$ and $\phi^{\varepsilon\ BPS}_{\text{net. inc.}}$ are based on approximations that neglect precautionary behavior. In particular, the expressions of the elasticities to shocks on net income as shares of lifetime expected income and current income in lifetime expected resources are derived from the same approximation of log-consumption growth as used in BPP, and they coincide with the expressions that would hold under perfect foresight, in the absence of uncertainty.[29] The exact expression of the elasticity of consumption to a permanent shock on net income is in fact the share of lifetime expected income minus the response of lifetime expected precautionary con-

[28]Though BPP note that approximated expression of the elasticities to transitory and permanent shocks on net income in the standard life-cycle model are the shares of current income and lifetime expected income in lifetime expected resources (bottom of p.1897), they do not use these expressions for estimation.

[29]In these two cases, consumption is a constant share of lifetime expected resources. A one unit permanent shock raises lifetime expected resources by an amount equal to lifetime expected income and current income, so its share in lifetime expected resources is the percentage change in consumption caused by a permanent shock when consumption is a constant fraction of lifetime expected resources. Similarly, a one unit transitory shock raises lifetime expected resources by an amount equal to current income.

sumption growth to a permanent income shock in lifetime expected resources net of lifetime expected precautionary consumption growth. The exact expression of the elasticity of consumption to a transitory shock on net income is in fact the share of current minus the response of lifetime expected precautionary consumption growth to a transitory income shock in lifetime expected resources net of lifetime expected precautionary consumption growth:

$$\phi^{\eta}_{net.inc.} = \frac{\sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s} - \overbrace{(dPG_{i,t}/d\eta_{i,t})}^{\neq 0 \text{ (precaution)}}}{(1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s} - \underbrace{PG_{i,t}}_{> 0 \text{ (precaution)}}} \neq \phi^{\eta\ BPS}_{net.inc.}$$

$$\phi^{\varepsilon}_{net.inc.} = \frac{y_t - \overbrace{(dPG_{i,t}/d\varepsilon_{i,t})}^{< 0 \text{ (precaution)}}}{(1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s} - \underbrace{PG_{i,t}}_{> 0 \text{ (precaution)}}} > \phi^{\varepsilon\ BPS}_{net.inc.},$$

with $PG_{i,t} = \sum_{s=1}^{T-t} \frac{l_{t,s} E_t[\varphi_{t+s-1}]}{(1+r)^s}$ the lifetime expected precautionary consumption growth, and $l_{t,s}$ a deterministic weight. I detail the derivation of these expressions in section E of the online appendix. Neglecting precautionary behavior induces an estimation bias in the measure of the Frisch elasticities, as it leads to plugging in expression $\phi^{\eta\ BPS}_{\text{net. inc.}}$ that wrongly estimates the sensitivity of consumption to a permanent shock, and expression $\phi^{\varepsilon\ BPS}_{\text{net. inc.}}$ that underestimates its sensitivity to a transitory shock.

Blundell, Pistaferri, and Saporta-Eksten (2016) develop a robustness check in which they allow $\phi^{\eta\ BPS}_{\text{net. inc.}}$, the share of expected lifetime income in expected lifetime resources, to be weighted by a parameter β that the authors estimate and find to be non significant. Yet, such a test would not necessarily capture the bias caused by the difference between $\phi^{\eta\ BPS}_{\text{net. inc.}}$ and its exact expression $\phi^{\eta}_{\text{net. inc.}}$, as their difference is not necessarily a constant fraction of $\phi^{\eta}_{\text{net. inc.}}$. Also, despite the presence of the parameter β, the model still imposes that the elasticities to a transitory shock be zero, and imposes a functional form $f(.)$ on $\phi^{\eta\ BPS}_{\text{wage}}$ that neglects precautionary behavior.

4.3 Estimation on biennial data

Method Since 1999, the PSID incorporates more comprehensive data on consumption but is only conducted every other year.[30] A number of studies apply the method de-

[30] Another biennial dataset that is sometimes used is the Bank of Italy Survey of Household Income and Wealth, which provides information on subjective income expectations in addition to consumption and income, but the panel period is also two years.

veloped by BPP to this more recent biennial panel data, changing the duration of the period from one year to two years. When the true frequency of income shocks is annual, though, using biennial data generates a bias. I denote Δ_2 the growth of a variable over two years. Log-income growth over two years is:

$$\Delta_2 ln(\tilde{y}_{i,t}) = \underbrace{\eta_{i,t} + \eta_{i,t-1}}_{\eta_{2i,t}} + \underbrace{\varepsilon_{i,t} + \theta\varepsilon_{i,t-1}}_{\varepsilon_{2i,t}} + \underbrace{\varepsilon_{i,t-2} + \theta\varepsilon_{i,t-3}}_{\varepsilon_{2i,t-2}},$$

where $\eta_{2i,t}$ is the permanent component of income when a period is two years, and $\varepsilon_{2i,t}$ its transitory component. When $\varepsilon_{i,t}$ is an MA(1) process, $\varepsilon_{2i,t}$ is an MA(0) process. The biennial version of the robust estimator ϕ_2^{ε} is:

$$\begin{aligned}
\phi^{\varepsilon}_{2\,i,t} &= \frac{cov(\Delta_2 ln(\tilde{c}_{i,t}), \varepsilon_{i,t} + \theta\varepsilon_{i,t-1})}{var(\varepsilon_{i,t} + \theta\varepsilon_{i,t-1})} \\
&= \frac{cov(\Delta_2 ln(\tilde{c}_{i,t}), \varepsilon_{i,t})}{var(\varepsilon_{i,t})} \times \underbrace{\frac{var(\varepsilon_{i,t})}{var(\varepsilon_{i,t}) + \theta^2 var(\varepsilon_{i,t-1})}}_{<1} + \theta \underbrace{\frac{cov(\Delta_2 ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1})}{var(\varepsilon_{i,t}) + \theta^2 var(\varepsilon_{i,t-1})}}_{\neq 0}.
\end{aligned}$$

Bias Although the instrument $\Delta_2 ln(\tilde{y}_{i,t+2})$ does not correlate with shocks prior to the realizations of $\Delta_2 ln(\tilde{c}_{i,t})$, it correlates with the shock $\varepsilon_{i,t-1}$ that occurs prior to $\varepsilon_{i,t}$, in between the two years of the survey. As a result, the covariance between log-consumption growth over two years and the instrument incorporates both the covariance between log-consumption growth over two years and the current transitory shock and the covariance between log-consumption growth over two years and the past transitory shock. The latter is non-zero because it includes the effect of the past transitory shock on current log-consumption growth and the effect of the past transitory shock on past log-consumption growth. Also, the covariance between log-income growth over two years and the instrument captures both the variance of the current transitory shock and a fraction θ^2 of the variance of the past transitory shock. There is therefore a bias, although its direction is undetermined, and the presence of a correlation between log-consumption growth and past shocks plays a role in the bias but is not its sole driver. Note that, if shocks are in fact drawn more frequently than every year, using yearly panel data as I do generates an estimation bias as well.

Table 5: Elasticity ϕ^{ε} and MPC lower bound $\underline{MPC}^{\varepsilon}$ with yearly and biennial estimators

Observation frequency	Yearly	Biennial
ϕ^{ε}	0.301	0.121
	(0.294)	(0.069)
$\underline{MPC}^{\varepsilon}$	0.166	0.067
	(0.162)	(0.038)
Obs.	3,694	3,694

Note: Standard errors in parentheses are adjusted for arbitrary within-household correlations and heteroskedasticity. Odd years are dropped from the sample. The first line reports GMM estimates of the average elasticity of nondurable consumption to a transitory shock on net income, computed with a moment restriction that uses yearly growth, and with a moment restriction that uses biennial growth. The second line reports a lower bound for the MPC out of the change in current income caused by a transitory shock. The third line reports the number of household-years for which the estimating moment is observed.

Quantitative importance In Table 5, I compare the yearly version of the robust estimator ϕ^{ε} to its biennial version ϕ_2^{ε}, to get a sense of the extent to which the use of biennial data shifts the elasticity estimate. I compute the value of one-year and two-year growth of current log-consumption, current log-income and future log-income for each household-year observation, and I drop every odd year starting in 1979, so that the shocks captured by the two estimators are the same. Results suggest that using a biennial estimator instead of a yearly estimator induces a downward bias, although the difference is not statistically significant because the elasticities are less precisely estimated with this substantially reduced subsample: the point estimate of the yearly estimator is 0.30, while the point estimate of the biennial estimator is 0.12, which is only 40% of the yearly point estimate. This is consistent with the presence of a negative effect of past transitory shocks on log-consumption growth pushing in the direction of a downwards bias.

5 Conclusion

In this paper, I show that the standard life-cycle model features a correlation between consumption growth and past transitory shocks that is caused by precautionary behavior, and that more general models possibly incorporate additional sources of correlation. I take stock of this possible correlation and generalize the semi-structural estimator of the elasticity of consumption to a transitory income shock, making it robust to such a correlation. The average elasticity of consumption to a transitory shock becomes statistically significant and its point estimate increases to 0.54, which is ten times larger than with an estimator that relies on instruments that are endogenous to past income shocks. The average marginal propensity to consume out of the total net present value change in income caused by a transitory shock is at least 0.24, consistent with the results obtained in natural experiments of transitory income changes.

What does it imply? First, the consistency of findings between the semi-structural estimation and the natural experiments suggests that the shocks considered in natural experiments are not too different from the typical shocks captured in longitudinal survey data. Thus, the strong response of consumption to a transitory income shock seems to be a widespread phenomena rather than a finding confined to fiscal stimuli and lottery wins. Second, the magnitude of the change in results when shifting from a non-robust to a robust estimator means that the effect of past shocks is not negligible and implies some caution in the use of other non-robust semi-structural techniques, including estimators of the elasticity to a permanent shock, estimators of Frisch elasticities, and estimators implemented in biennial datasets.

Appendix A Proof of the theorem

Theorem: In the model presented in the first section, the precautionary consumption growth, φ_t, is negatively correlated with assets. At any period $0 < t < T$:

$$\frac{d\varphi_{i,t}}{da_{i,t}} < 0.$$

Proof: A household solves the maximization problem described by (2.1)-(2.4). To see how a change in assets affects its allocation of consumption over time, I derive both sides of the first order condition of the maximization problem with respect to a change in assets, and divide by $(-u''(c_t R_{i,t,t+1}^{1/\rho}))$:

$$\frac{dc_{i,t}R_{i,t,t+1}^{1/\rho}}{da_{i,t}} = E_t[\frac{dc_{i,t+1}}{da_{i,t}}\frac{u''(c_{i,t+1})}{u''(c_{i,t}R_{i,t,t+1}^{1/\rho})}]$$

$$\frac{dc_{i,t}R_{i,t,t+1}^{1/\rho}}{da_{i,t}} = \frac{dE_t[c_{i,t+1}]}{da_{i,t}}\underbrace{\frac{E_t[-u''(c_{i,t+1})]}{-u''(c_{i,t}R_{i,t,t+1}^{1/\rho})}}_{>1} + \underbrace{cov_t\left(\frac{dc_{i,t+1}}{da_{i,t}}, \frac{-u''(c_{i,t+1})}{-u''(c_{i,t}R_{i,t,t+1}^{1/\rho})}\right)}_{>0}.$$

Two effects lead a household to raise its current consumption more than its expected future consumption in response to a gain in assets, reducing the precautionary gap between its current and expected future consumption. First, when utility is isoelastic, the change in marginal utility $-u''(.)$ is a convex function of marginal utility $u'(.)$. At the point where the distributions of current and future consumption equalize current and future expected marginal utility, they do not equalize current and future expected change in marginal utility, since the change in marginal utility is more convex: $-u''(c_{i,t}) < E_t[-u''(c_{i,t+1})]$.[31] Thus, if current and future consumption increased by the same amount in response to a gain in assets, the marginal utility of current consumption would fall less than the marginal utility of future consumption is expected to. For them to fall by the same amount, current consumption must increase more than future consumption. Intuitively, when the change in marginal utility $-u''(.)$ is a convex function of marginal utility $u'(.)$, the convexity of the marginal utility is less pronounced over higher levels of future consumption: a shift upwards in future consumption by d is equivalent to a change in marginal utility from $u'(.)$ to the less convex function $u'(.) - d(-u''(.))$. As a result, the expected marginal utility of future consumption falls more than the marginal utility of current consumption when consumption shifts up, because of the reduced effect of uncertainty caused by the reduced convexity.

Second, a change in assets does not raise future consumption by the same amount $\frac{dE_t[c_{i,t+1}]}{da_{i,t}}$ in all states of the world. When utility is isoelastic, the cross-derivatives of consumption with respect to a change in assets and to a transitory or a permanent shock are negative.[32] This means that a gain in assets causes a larger increase in future consumption if an adverse income shock realizes, which is when the marginal utility of future consumption is high and falls a lot with an increase in consumption, than if a favorable income shock realizes, which is when the marginal utility of future consumption is low and does not fall a lot with an increase in consumption. The covariance between the response of future consumption to a change in assets and the change in the marginal utility of future consumption is positive: $cov_t\left(\frac{dc_{i,t+1}}{da_{i,t}}, \frac{-u''(c_{i,t+1})}{-u''(c_{i,t}R_{i,t,t+1}^{1/\rho}}\right) > 0$. The response of consumption is higher in the states of the world when an increase in consumption reduces the marginal utility of consumption most, so the expected marginal utility of future consumption falls more with a gain in assets than it would if future consumption responded in the same way in all states of the world. Current consumption must increase even more than expected future consumption in response to a gain in assets for its marginal utility to fall as much. Intuitively, because a gain in assets raises

[31] From Arrow-Pratt (1964), since $-u''(.)$ is a convex function of $u'(.)$, the risk premium $\varphi_{i,t}^{-u''}$ associated with $-u''(.)$ is larger than the risk premium $\varphi_{i,t} = \varphi_{i,t}^{u'}$ associated with $u'(.)$, and: $E_t[-u''(c_{i,t+1})] = -u''(E_t[c_{i,t+1}] - \varphi_{i,t}^{-u''}]) > -u''(E_t[c_{i,t+1}] - \varphi_{i,t}^{u'}]) = -u''(c_{i,t})$

[32] These results are proved in Commault (2018b).

consumption most in the states of the world when it is lowest, it reduces the variance of future consumption, reducing the effect of uncertainty on the expected marginal utility even further.

Appendix B Consistency of the estimator

I can consistently detrend log-consumption from the effect of current demographic characteristics because current characteristics are independent of the current and past income shocks, which form the residuals of the detrending regressions. The first difference of this residual is:

$$\Delta ln(\tilde{c}_{i,t}) = f(\varepsilon_{i,t}, \eta_{i,t}, \zeta^c_{i,t}\varepsilon_{i,t-1}, \eta_{i,t-1}, \zeta^c_{i,t-1}, ...).$$

I take an exact Taylor expansion of (residual) log-consumption growth around the point where the realization of the current transitory shock is equal to zero:

$$\Delta ln(\tilde{c}_{i,t}) = f(0, \eta_{i,t}, \zeta^c_{i,t}\varepsilon_{i,t-1}, \eta_{i,t-1}, \zeta^c_{i,t-1}, ...) + \sum_{s=1}^{\infty} \frac{\varepsilon^s_{i,t}}{s!} \left(\frac{d^s \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^s_{i,t}} \right)_{|0},$$

where the subscript $|_0$ indicates that the variable is considered at the point where $\varepsilon_{i,t} = 0$. The elasticity of consumption to a transitory shock of household i at period t is:

$$\phi^{\varepsilon}_{i,t} = \frac{d\Delta ln(\tilde{c}_{i,t})}{d\varepsilon_{i,t}} = \sum_{s=1}^{\infty} \frac{\varepsilon^{s-1}_{i,t}}{(s-1)!} \left(\frac{d^s \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^s_{i,t}} \right)_{|0}.$$

It writes as a polynomial of the current transitory shock. I denote $\phi^{\varepsilon} = E[\phi^{\varepsilon}_{i,t}]$ the average value of the elasticity in the sample. As a transitory shock is independent of the permanent shocks, of the other shocks, and of its own past realizations, it is independent of $\left(\frac{d^s \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^s_{i,t}} \right)_{|0}$. The average value of the product of the transitory shock and the derivatives of log-consumption growth at the point where $\varepsilon_{i,t} = 0$ is the product of their average values:

$$E[\phi^{\varepsilon}_{i,t}] = \sum_{s=1}^{\infty} \frac{E[\varepsilon^{s-1}_{i,t}]}{(s-1)!} E\left[\left(\frac{d^s \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^s_{i,t}} \right)_{|0} \right].$$

Because each household draws its transitory shock from a normal distribution, and because the moment m of a variable x that is normally distributed is $E[x^m] = \mathbb{1}_{\{m \text{ is even}\}} E[x^2]^{m/2}(m-$

1)!!, the expression is:[33]

$$E[\phi^{\varepsilon}_{i,t}] = \sum_{s=1}^{\infty} \mathbb{1}_{\{(s-1) \text{ is even}\}} E[\varepsilon^2_{i,t}]^{(s-1)/2} \frac{(s-2)!!}{(s-1)!} \left(\frac{d^s \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^s_{i,t}} \right)_{|0}$$

$$= \sum_{s=1}^{\infty} \mathbb{1}_{\{(s-1) \text{ is even}\}} \frac{E[\varepsilon^2_{i,t}]^{(s-1)/2}}{(s-1)!!} E\left[\left(\frac{d^s \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^s_{i,t}} \right)_{|0} \right].$$

A consistent estimator of this average elasticity is the ratio of the covariance between log-income growth and the transitory shock over the variance of the transitory shock:

$$\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})}{var(\varepsilon_{i,t})} = \frac{1}{E[\varepsilon^2_{i,t}]} \sum_{s=1}^{\infty} \frac{E[\varepsilon^{s+1}_{i,t}]}{(s)!} E\left[\left(\frac{d^s \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^s_{i,t}} \right)_{|0} \right]$$

$$= \sum_{s=1}^{\infty} \mathbb{1}_{\{(s+1) \text{ is even}\}} \frac{1}{E[\varepsilon^2_{i,t}]} E[\varepsilon^2_{i,t}]^{(s+1)/2} \frac{s!!}{s!} E\left[\left(\frac{d^s \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^s_{i,t}} \right)_{|0} \right]$$

$$= E[\phi^{\varepsilon}_{i,t}].$$

Assuming that the moments of order higher than two are zero, and that $\frac{d^3 \Delta ln(\tilde{c}_{i,t})}{d\varepsilon^3_{i,t}} = 0$ would also make the ratio $\frac{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})}{var(\varepsilon_{i,t})}$ a consistent estimator of the elasticity. An alternative hypothesis under which this estimator is straightforwardly consistent is when the effect of a contemporaneous transitory shock on log-consumption growth is linear and additively separable from the effect of other shocks.

Appendix C Data

The main data source is the PSID, which contains longitudinal information on a representative sample of US households, surveyed every year. It started in 1968 with approximately 3,000 households. Both the original households and their splitoffs have been followed since. The period I consider is 1978-1992.[34] I select out households that are not continuously married over the period, those experiencing a dramatic change in family composition, those headed by a female, those with missing reports on race, education, and region, and those whose head is younger than 30 or older than 65. I also drop

[33] When households are not drawing shocks all from the same normal distribution, but from J different normal distributions, it writes: $E[x^m] = E[x^m | j = j_1] + .. + E[x^m | j = j_J] = \mathbb{1}_{\{m \text{ is even}\}} E[x^2]^{m/2}(m-1)!!$. The only caveat is that I need a sufficiently large number of households drawing from each distribution j for the sample averages to converge towards their theoretical expressions.

[34] The CEX data that is used to impute consumption is difficult to use before 1978. After 1992, a number of the questions used by BPP to build their measure of income are redesigned.

some income outliers. The dataset, the period, and the selection are the same as in BPP. The final sample is composed of 15,779 household-year observations from 1,765 households[35]. Among these, there are 12,041 household-year observations for which current log-consumption growth and current log-income growth are simultaneously observed, and 8,958 for which current log-consumption growth, current log-income growth, and log-income growth two periods later are simultaneously observed.

Net income is the taxable family income reported by a household minus its financial income and minus the federal taxes paid on nonfinancial income.[36] Gross income is net income plus taxes. Gross income before transfers is gross income minus transfer income.[37] All three measures are deflated by the contemporaneous Consumer Price Index (CPI).

Nondurable consumption is the sum of annual expenditure on food, alcohol, tobacco, nondurable services, heating fuel, public and private transport (including gasoline), personal care, and clothing and footwear, deflated by the CPI. Total consumption is the sum of nondurable consumption plus annual expenditures on durable goods, namely housing (mortgage interest, property tax, rent, other lodging, textiles, furniture, floor coverings, appliances), new and used cars, vehicle finance charges and insurance, car rentals and leases, cash contributions, and personal insurance (life insurance and retirement), deflated by the CPI. Total consumption plus health and education is the sum of total consumption plus annual expenditures on health (insurance, prescription drugs, medical services), and education. As the PSID only reports expenditure on food, these three measures of consumption are imputed from the demographic characteristics of the households and from their food consumption, with the coefficients used for the imputation estimated from the CEX over the same period. Further details are provided in the paper of BPP (section I.B.).

I detrend log-income and log-consumption from the impact of demographic variables by regressing them on dummies for year, year-of-birth, family size, number of children, existence of outside dependent children, education, race, employment status, presence of an additional income recipient that is not the head or his spouse, region, residence in a large city, and interactions between a subset of these demographic charac-

[35]My sample is exactly the same as that of BPP. The number of household-year observations reported differ (they report 17,604 observations) simply because they count observations of log-income, while I report observations of log-income growth, which is the variable used for estimation.

[36]Federal taxes on nonfinancial income are assumed to be a proportion of total federal taxes; the proportionality coefficient is given by the ratio of nonfinancial income over total income.

[37]Transfer income includes aid to families with dependent children, supplemental security income and other welfare payments, social security income and other retirement, pensions and annuities payments, unemployment benefits, worker's compensations, child support, help from relatives, and other transfer income.

teristics and year and cohort dummies. This follows BPP, except that I add interactions with cohort dummies, which are present in a more recent version of the BPP estimator (the one used by Blundell, Pistaferri, and Saporta-Eksten (2016)).

Appendix D Estimating with a generalized method of moment

I estimate the average elasticity with a generalized method of moments. The statistical model implies that:

$$cov(\Delta ln(\tilde{c}_{i,t}), \Delta ln(\tilde{y}_{i,t+2})) = E[\Delta ln(\tilde{c}_{i,t}) \Delta ln(\tilde{y}_{i,t+2})] = \phi^{\varepsilon} var(\varepsilon_{i,t}),$$
$$cov(\Delta ln(\tilde{y}_{i,t}), \Delta ln(\tilde{y}_{i,t+2})) = E[\Delta ln(\tilde{y}_{i,t}) \Delta ln(\tilde{y}_{i,t+2})] = var(\varepsilon_{i,t}).$$

Thus, the following moment restriction holds:

$$E[\underbrace{\Delta ln(\tilde{c}_{i,t}) \Delta ln(\tilde{y}_{i,t+2}) - \phi^{\varepsilon} \Delta ln(\tilde{y}_{i,t}) \Delta ln(\tilde{y}_{i,t+2})}_{g(X_{i,t}, \phi^{\varepsilon})}] = 0,$$

with $X_{i,t} = (\Delta ln(\tilde{c}_{i,t}), \Delta ln(\tilde{y}_{i,t}), \Delta ln(\tilde{y}_{i,t+2}))$ the set of variables involved , and $\phi^{\varepsilon} = E[\phi^{\varepsilon}_{i,t}]$ the parameter involved. This restriction makes it possible to estimate the parameter ϕ^{ε} as the value that minimizes a norm of the sample analog of this moment:

$$\hat{\phi^{\varepsilon}} = \underset{\phi^{\varepsilon}}{\operatorname{argmin}} \left(\frac{1}{N} \sum_{n=1}^{N} g(X_n, \phi^{\varepsilon}) \right)^{\mathsf{T}} \hat{W} \left(\frac{1}{N} \sum_{n=1}^{N} g(X_n, \phi^{\varepsilon}) \right),$$

with N the number of household-year observations (i,t) at which the three variables $\Delta ln(\tilde{c}_{i,t})$, $\Delta ln(\tilde{y}_{i,t})$, and $\Delta ln(\tilde{y}_{i,t+2})$ are observed, and $\hat{W}$ a weighting matrix. The matrix is chosen so the estimation of the standard error is robust to arbitrary within-household correlations and robust to heteroskedasticity, that is, to the fact that the residuals $g(X_{i,t}, \phi^{\varepsilon}) - E[g(X_{i,t}, \phi^{\varepsilon})]$ are not drawn from the same distribution.

References

Agarwal, S., Liu, C., and Souleles, N. (2007). "The Reaction of Consumer Spending and Debt to Tax Rebates-Evidence from Consumer Credit Data." *Journal of Political Economy* 115.6, pp. 986–1019.

Agarwal, S. and Qian, W. (2014). "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore." *American Economic Review* 104.12, pp. 4205–30.

Arellano, M., Blundell, R., and Bonhomme, S. (2017). "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework." *Econometrica* 85, pp. 693–734.

Baker, S. and Yannelis, C. (2017). "Income Changes and Consumption: Evidence from the 2013 Federal Government Shutdown." *Review of Economic Dynamics* 23, pp. 99–124.

Blundell, R., Low, H., and Preston, I. (2013). "Decomposing changes in income risk using consumption data." *Quantitative Economics* 4.1, pp. 1–37.

Blundell, R., Pistaferri, L., and Preston, I. (2008). "Consumption Inequality and Partial Insurance." *American Economic Review* 98.5, pp. 1887–1921.

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). "Consumption inequality and family labor supply." *American Economic Review* 106.2, pp. 387–435.

Carroll, C. and Kimball, M. (1996). "On the Concavity of the Consumption Function." *Econometrica* 64.4, pp. 981–92.

Cochrane, J. (1991). "A Simple Test of Consumption Insurance." *Journal of Political Economy* 99.5, pp. 957–76.

Commault, J. (2018b). "The Concavity of the Consumption Function Revisited." mimeo.

Deaton, Angus (1991). "Saving and Liquidity Constraints." *Econometrica* 59.5, pp. 1221–1248.

Drèze, J. and Modigliani, F. (1972). "Consumption decisions under uncertainty." *Journal of Economic Theory* 5.3, pp. 308–335.

Dynarski, S. and Gruber, J. (1997). "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity* 28.1, pp. 229–303.

Fagereng, A., Holm, M., and Natvik, G. (2018). "MPC Heterogeneity and Household Balance Sheets." mimeo.

Gelman, M., Kariv, S, Shapiro, M., Silverman, D., and Tadelis, S (2018). "How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown." *Journal of Public Economics*.

Hall, R. (1978). "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86.6, pp. 971–987.

Johnson, D., Parker, J., and Souleles, N. (2006). "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96.5, pp. 1589–1610.

Kan, K., Peng, S., and Wang, P. (2017). "Understanding Consumption Behavior: Evidence from Consumers' Reaction to Shopping Vouchers." *American Economic Journal: Economic Policy* 9.1, pp. 137–53.

Kaplan, G. and Violante, G. (2010). "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics* 2.4, pp. 53–87.

Kaplan, G. and Violante, G. (2014). "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82.4, pp. 1199–1239.

Kimball, M. (1990). "Precautionary Saving in the Small and in the Large." *Econometrica* 58.1, pp. 53–73.

Krueger, D. and Perri, F. (2005). "Understanding Consumption Smoothing: Evidence from the U.S. Consumer Expenditure Data." *Journal of the European Economic Association* 3.2-3, pp. 340–349.

Krueger, D. and Perri, F. (2011). "How do Households Respond to Income Shocks?" mimeo.

Leland, H. (1968). "Saving and Uncertainty: The Precautionary Demand for Saving." *The Quarterly Journal of Economics* 82.3, pp. 465–473.

Meghir, C. and Pistaferri, L. (2004). "Income Variance Dynamics and Heterogeneity." *Econometrica* 72.1, pp. 1–32.

Miller, B. (1976). "The effect on optimal consumption of increased uncertainty in labor income in the multiperiod case." *Journal of Economic Theory* 13.1, pp. 154 –167. ISSN: 0022-0531.

Misra, K. and Surico, P. (2014). "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs." *American Economic Journal: Macroeconomics* 6.4, pp. 84–106.

Parker, J. (1999). "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes." *The American Economic Review* 89.4, pp. 959–973.

Parker, J. and Souleles, N. (2017). "Reported Preference vs. Revealed Preference: Evidence from the Propensity to Spend Tax Rebates." mimeo.

Parker, J., Souleles, N., Johnson, D., and McClelland, R. (2013). "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103.6, pp. 2530–53.

Sandmo, A. (1970). "The Effect of Uncertainty on Saving Decisions." *Review of Economic Studies* 37.3, pp. 353–360.

Sibley, D. (1975). "Permanent and transitory income effects in a model of optimal consumption with wage income uncertainty." *Journal of Economic Theory* 11.1, pp. 68–82.

Souleles, N. (1999). "The Response of Household Consumption to Income Tax Refunds." *American Economic Review* 89.4, pp. 947–958.

Souleles, N. (2002). "Consumer response to the Reagan tax cuts." *Journal of Public Economics* 85.1, pp. 99–120.